

Operations Guide

hp StorageWorks Continuous Access EVA V1.1B

Product Version: 1.1B

Sixth Edition (July 2004)

Part Number: AA-RTEHF-TE

HP StorageWorks Continuous Access EVA is a Fibre Channel storage solution that uses controller-based replication to provide disaster-tolerant data through the use of hardware redundancy across several sites. This document provides instructions for configuring and operating this solution for the Enterprise Virtual Array.



© Copyright 2003–2004 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

Java™ is a U.S. trademark of Sun Microsystems, Inc.

Microsoft®, Windows®, and Windows NT® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided “as is” without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements for such products. Nothing herein should be construed as constituting an additional warranty.

Printed in the U.S.A.

Continuous Access EVA V1.1B Operations Guide
Sixth Edition (July 2004)
Part Number: AA–RTEHF–TE

Contents

About this Guide	11
Revision history	12
Overview	13
Intended audience	13
Prerequisites	13
Related documentation	13
Conventions	14
Document conventions	14
Text symbols	15
Getting help	15
HP technical support	15
HP storage website	16
HP authorized reseller	16
1 About Continuous Access EVA	17
Overview of Continuous Access EVA	17
Hardware components	19
Enterprise Virtual Array	19
Fibre Channel SAN switches	20
Storage Management Appliance	21
Fibre Channel adapters	21
Required software	22
Virtual Controller Software	22
Storage Management Appliance software	22
Host operating systems	24
Secure Path	24
Licensing	25
Business Copy EVA license	25
Continuous Access EVA license	25

2	Concepts	27
	Virtualization concepts	28
	Physical vs. virtual storage	28
	Benefits over traditional storage	28
	Virtual RAID types	29
	Business Copy concepts	31
	Snapshots	31
	Snapclones	32
	Remote mirrors	32
	Continuous Access EVA concepts	33
	Remote data replication	33
	Copy sets	34
	DR groups	34
	DR group properties	36
	Log disks	37
	Managed sets	38
	Failover	39
	Zoning	41
3	Configuration Planning	43
	About disk groups	44
	Planning your disk group configuration	46
	Naming restrictions	48
	Planning fabric configurations	49
	Dual-fabric configuration	49
	Single-fabric configurations	50
	Load balancing	52
	Planning your zones	53
	Restrictions	56
4	Configuration	59
	Hardware configuration	60
	Fibre Channel adapter installation	60
	Configuring the Fibre Channel switches	61
	Controller-to-switch connections	61
	Host-to-switch connections	62
	EVA zoning recommendations	62
	Zoning with B-series switches	62
	Zoning with C- and M-series fabric switches	64

HSG80 zoning recommendations	64
SMA-to-switch connections	65
Software configuration	65
Storage Management Appliance software setup	66
Command View EVA	66
Continuous Access user interface	66
Recommended uses for the Continuous Access user interface	66
Monitoring with the Continuous Access user interface	67
Configuring hosts	68
Licensing	69
Initializing your storage systems	69
Creating disk groups	69
Creating host folders	70
Creating hosts	70
Creating Vdisk folders	70
Creating Vdisks	70
Installing Secure Path	72
Presenting Vdisks to hosts	72
Accessing the Continuous Access user interface	72
Creating copy sets and DR groups	75
Presenting a copy set to a destination host	79
Specifying disk group membership for a log	79
Deleting or detaching copy sets	81
Deleting DR groups	81
Creating managed sets	81
Editing a managed set	82
Adding a DR group to a managed set	82
Removing a DR group from a managed set	83
Deleting a managed set	83
Backing up configuration information	84
5 Storage Management Appliance Procedures	87
Considerations for managing storage with multiple SMAs	88
Saving your Continuous Access EVA storage configuration	89
Stopping SMA applications	90
Restarting SMA applications	92
Moving storage management to another SMA	93
Synchronizing time on the SMAs	96
Setting a storage system time to an SMA	96

Synchronizing SMAs to an NTP server	98
Enabling management on an SMA when HSG80 controllers are present	99
The managing SMA is still accessible	99
The managing SMA is disabled	100
6 Recovery	103
Planning for a disaster	104
Failsafe and normal modes	105
Throttling of merge I/O after logging	106
Failover defined	106
The Continuous Access user interface	110
Continuous Access user interface icons	110
Data replication using the Continuous Access user interface	113
Suspend	114
Resume	114
Failover	114
Disable Failsafe	115
Possible event scenarios	115
Planned failover	116
Unplanned failover	116
Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)	117
Return operations to home storage system	117
Return operations to replaced new storage hardware	117
Disk group hardware failure on the source storage system	118
Disk group hardware failure on the destination storage system	118
Failover and recovery procedures	119
Planned failover	119
Unplanned failover	127
Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)	132
Return operations to home storage system	134
Return operations to replaced new storage hardware	134
Recovering from a disk group hardware failure	142
Disk group hardware failure on the source storage system	143
Recovery when data replication was normal before failure	143
Recovery when data replication was logging before failure	149
Disk group hardware failure on the destination storage system	150

7	Troubleshooting with Multiple Sites	153
	Troubleshooting storage problems	153
	Troubleshooting intersite link problems	175
8	Continuous Access EVA Support Procedures	177
	Creating a destination Snapclone before making a full copy	177
	Data movement using a Snapclone	178
	Three-site cascaded data replication using Snapclones	180
	Bootless DR group planned failover with Linux using LVM in standalone mode or with SuSE SLES 8 running LifeKeeper 4.4.3	197
	Source host procedure	197
	Destination host procedure	198
	Performing a failover back to the previous source	199
A	Description of Event and Termination Codes	201
	Glossary	219
	Index	233
	Figures	
1	Basic Continuous Access EVA configuration with redundant Storage Management Appliances	18
2	A typical EVA rack configuration (Model 2C12D)	20
3	A typical Fibre Channel SAN switch	21
4	Functional relationship between controllers and SMA	23
5	Traditional RAID vs. virtual RAID	30
6	Continuous Access EVA DR group replication	35
7	Managed sets	38
8	Replicating relationships among DR groups	40
9	Simple Continuous Access EVA zoning	41
10	Building sequential disk groups	45
11	Dual-fabric configuration	49
12	A single-switch fabric configuration	50
13	Two-switch, single-fabric configuration	51
14	Mesh-switch fabric configuration	51
15	SMA management zones	54
16	Controller-to-switch cabling	62
17	Example of host zoning with infrastructure switches	63

18	HSG Element Manager Controller Properties page.....	65
19	Event notification flow.....	68
20	SMA Device Home page.....	73
21	SMA login window.....	73
22	SMA software Home page.....	74
23	SMA Tools page.....	74
24	Continuous Access user interface main window.....	75
25	Storage selection on Continuous Access user interface main window.....	76
26	Add the new copy set window.....	76
27	Create a new DR Group window.....	77
28	Create a new Copy Set window.....	78
29	Storage allocation check on Vdisk Active Member Properties page.....	79
30	Edit or Create a Managed Set window.....	82
31	Continuous Access user interface Select Location window.....	89
32	Continuous Access user interface Download window.....	90
33	SMA Settings page.....	91
34	SMA Manage Tools page.....	91
35	Selected applications to stop on the Manage Tools page.....	92
36	Selected applications to restart on the Manage Tools page.....	93
37	Storage System Managed by Another Agent page.....	94
38	Assuming management of your storage message.....	94
39	Initialized Storage System Properties page.....	96
40	System Options page.....	97
41	Set System Time page.....	97
42	Date and Time Setting page.....	98
43	Data Replication menu on the Continuous Access user interface.....	113
44	Planned and unplanned transfer of operations.....	120
45	Resumption of operations if unable to access destination in failsafe mode.....	133
46	Return operations to new hardware.....	135
47	Normal disk group indication with Continuous Access user interface.....	144
48	Hardware failure viewed from Continuous Access user interface.....	145
49	Hardware failure viewed from Command View EVA.....	145
50	Command View EVA hardware deletion process.....	146
51	Message confirming Vdisk and DR group deletion.....	147
52	Command View EVA manual deletion message.....	147
53	Message confirming DR group deletion.....	148
54	Vdisk Deletion in Progress message.....	148
55	DR group Properties page.....	150

56	DR group deletion message	151
57	Disk group Hardware Failure page	151
58	Creating a DR group from a Snapclone	179
59	Data movement using Snapclones example	181
60	Setting a DR group to failsafe mode	182
61	Setting a DR group for synchronous replication	183
62	Checking normalization of DR group members	183
63	Host Folder Properties page	184
64	Add a Host page	185
65	Operation succeeded page	185
66	Vdisk Active Member Properties page	186
67	Create a Snapclone page 1	186
68	Vdisk Active Member Properties page with Add member tab	188
69	Create a DR Group page	189
70	Vdisk Active Member Properties page with Unpresent tab	190
71	Unpresent Host(s) page	190
72	Vdisk Active Member Properties with Remove member tab	191
73	Keep/Delete Remote Mirror Vdisk page	192
74	Vdisk Family Properties window	192
75	Changing Write Protect attribute of Vdisk	193
76	Vdisk Active Member Properties page with Present tab	194
77	Present Vdisk page	195
78	Vdisk Family Properties page	196

Tables

1	Revision history	12
2	Document Conventions	14
3	Continuous Access EVA Platform Zoning Requirements	53
4	Zoning input form	55
5	Restrictions	56
6	License key types	69
7	When and when not to fail over a DR group, managed set, or storage system	108
8	Continuous Access user interface icons	111
9	Identifying your troubleshooting situation	154
10	Troubleshooting a site storage system with single or multiple site replicating relationships	156
11	Continuous Access EVA event codes with Command View EVA descriptions	202

12	Continuous Access EVA termination codes and descriptions	213
13	Event codes with Continuous Access user interface summary description	217

About This Guide

This operations guide provides information to help you:

- Understand HP StorageWorks Continuous Access EVA hardware and configuration requirements
- Understand virtualized storage systems and Continuous Access EVA concepts
- Monitor controller events
- Perform failovers and recovery procedures
- Contact technical support for additional assistance

“About this Guide” topics include:

- [Revision history](#), page 12
- [Overview](#), page 13
- [Conventions](#), page 14
- [Getting help](#), page 15

Revision history

[Table 1](#) identifies significant changes made to this document during each revision.

Table 1: Revision history

Edition/Date	Change summary
First (May 2003)	Initial publication
Second (July 2003)	<ul style="list-style-type: none"> ■ Increased the number of EVAs and FCAs per replicating pair of EVAs in the configuration limits. ■ Added Microsoft® Windows® 2003 (32-bit) support. ■ Added single switch/fabric/FCA discussion to Configuration Planning chapter. ■ Renamed Infrastructure and High Availability switches to B-series and M-series switches, respectively. ■ Renamed “Failover” chapter to “Recovery” and revised “Return Operations to Replaced New Hardware” procedure in this chapter.
Third (October 2003)	<ul style="list-style-type: none"> ■ Added support for these VCS 3.010 features: <ul style="list-style-type: none"> — EVA3000 — 128 copy sets, 128 DR groups — asynchronous replication — multiple replication relationships ■ Added chapter on troubleshooting with multiple sites. ■ Added event and termination codes to appendix.
Fourth (December 2003)	<ul style="list-style-type: none"> ■ Increased maximum number of FCAs supported to 256 per EVA. ■ Removed support for C-series switches until further notice.
Fifth (December 2003)	Added support for C-series switches.
Sixth (July 2004)	Added support for: <ul style="list-style-type: none"> ■ VCS V3.02 and Command View EVA V3.2 ■ 300 GB online and 250 GB near-online FATA drives Added bootless DR group failover procedure for Linux.

Overview

This section covers the following topics:

- [Intended audience](#)
- [Prerequisites](#)
- [Related documentation](#)

Intended audience

This document is intended for system and network administrators who are experienced with the following:

- SAN fabric configurations
- Continuous Access EVA-supported host operating system environments
- Enterprise Virtual Array storage systems

Prerequisites

This document assumes the user has:

- Decided on a Continuous Access EVA design and ordered the components
- Configured dual fabrics with working intersite links

Related documentation

In addition to this document, HP provides the following related information. For the Continuous Access EVA documentation listed below, go to:

<http://h18006.www1.hp.com/products/storage/software/conaccesseva/index.html>

- *HP StorageWorks Continuous Access EVA V1.1 Getting Started Guide*
- *HP StorageWorks Continuous Access EVA V1.1B Design Reference Guide*
- *HP StorageWorks Continuous Access EVA V1.1B Release Notes*
- *HP StorageWorks Continuous Access User Interface V1.1A Installation Guide*
- *HP StorageWorks Continuous Access User Interface V1.1A Release Notes*
- *HP StorageWorks Continuous Access EVA and Data Replication Manager SAN Extensions Reference Guide*

For the EVA3000, go to:

<http://h18006.www1.hp.com/products/storageworks/eva3000/index.html>

For the EVA5000, go to:

<http://h18006.www1.hp.com/products/storageworks/enterprise/index.html>

For Storage Management Appliances or Command View EVA, go to:

<http://h18006.www1.hp.com/products/sanworks/managementappliance/index.html>

For SAN design or SAN extensions, go to:

<http://h18006.www1.hp.com/products/storageworks/san/documentation.html>

Conventions

Conventions consist of the following:

- [Document conventions](#)
- [Text symbols](#)

Document conventions

The document conventions included in [Table 2](#) apply to this document.

Table 2: Document Conventions

Element	Convention
Cross-reference links	Blue text: Figure 1
Key and field names, menu items, buttons, and dialog box titles	Bold
File names, application names, and text emphasis	<i>Italics</i>
User input, command and directory names, and system responses (output and messages)	Monospace font COMMAND NAMES are uppercase monospace font unless they are case sensitive
Variables	<monospace, italic font>
Website addresses	Blue, underlined sans serif font text: http://www.hp.com

Text symbols

The following symbols may be found in the text of this guide. They have the following meanings.



WARNING: Text set off in this manner indicates that failure to follow directions in the warning could result in bodily harm or death.



Caution: Text set off in this manner indicates that failure to follow directions could result in damage to equipment or data.

Note: Text set off in this manner presents commentary, sidelights, or interesting points of information.

Getting help

If you still have a question after reading this guide, contact an HP authorized service provider or access our website: <http://www.hp.com>.

HP technical support

Telephone numbers for worldwide technical support are listed on the following HP website: <http://www.hp.com/support/>. From this website, select the country of origin.

Note: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages

- Operating system type and revision level
- Detailed, specific questions

HP storage website

The HP website has the latest information on this product, as well as the latest drivers. Access storage at <http://www.hp.com/country/us/eng/prodserv/storage.html>. From this website, select the appropriate product or solution.

HP authorized reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.
- Elsewhere, see the HP website for locations and telephone numbers: <http://www.hp.com>.

About Continuous Access EVA

1

This chapter provides an overview of Continuous Access EVA features and gives a brief description of the hardware components and software applications that are required.

Topics in this chapter include:


- [Overview of Continuous Access EVA](#), page 17
- [Hardware components](#), page 19
 - [Enterprise Virtual Array](#), page 19
 - [Fibre Channel SAN switches](#), page 20
 - [Storage Management Appliance](#), page 21
 - [Fibre Channel adapters](#), page 21
- [Required software](#), page 22
 - [Virtual Controller Software](#), page 22
 - [Storage Management Appliance software](#), page 22
 - [Host operating systems](#), page 24
 - [Secure Path](#), page 24
- [Licensing](#), page 25
 - [Business Copy EVA license](#), page 25
 - [Continuous Access EVA license](#), page 25

Overview of Continuous Access EVA

Continuous Access EVA is a solution for mirroring data online and in real time to remote locations via a local or an extended storage area network (SAN). Using controller and management software, data replication is performed at the storage system level and in the background to any host activity.

Continuous Access EVA provides disaster-tolerant data through the use of hardware redundancy and data replication between two EVA storage systems that are separated by a safe distance. (For a discussion of safe distance, refer to the *HP StorageWorks Continuous Access EVA Design Reference Guide*.) Multiple hosts can be connected to one or more shared storage systems that run homogeneous or heterogeneous operating systems.

Figure 1 shows a basic Continuous Access EVA configuration consisting of two separated storage systems. One is located at a *primary* (or *local*) site and the other at an *alternate* (or *remote*) site. In this illustration, two redundant fabrics are shown, one called the *gray fabric* and the other called the *black fabric*. Each storage system can perform primary data processing functions as a *source*, with data replication occurring on the *destination* storage system. The replication process can also be *bidirectional*, with some I/O streams moving to the storage system and other I/O streams moving simultaneously from the storage system. This feature allows the storage system to be the source for some data groups and the destination for others.



CX08165B

Figure 1: Basic Continuous Access EVA configuration with redundant Storage Management Appliances

If a significant failure occurs at the source storage system location, hardware and software redundancy allows data processing to quickly resume at the destination storage system, where the data is intact. This process is called *failover*. When the cause of the storage system failure has been resolved, data processing can be moved back to the original source storage system by performing another failover.

The basic Continuous Access EVA configuration uses one or more HP OpenView Storage Management Appliances (SMAs) at each site. Multiple SMAs provide redundancy and allow storage systems to be individually managed by specific SMAs.

Hardware components

The following sections describe individual hardware components that are necessary to build a Continuous Access EVA solution. Depending on the size of your SAN and the considerations used in designing it, many different hardware configurations are possible. Refer to the *HP StorageWorks Continuous Access EVA Design Reference Guide* for a detailed description of various hardware configurations.

Enterprise Virtual Array

Continuous Access EVA uses a minimum of two EVA storage systems: one at a primary site and one at an alternate site. Each EVA storage system is housed in a rack like that shown in [Figure 2](#). This is one of many possible EVA rack configurations. Additional configurations are described in the *HP StorageWorks Enterprise Virtual Array Hardware Configuration Guide*.

The rack houses HSV controllers and as many as 12 drive shelves that make up a storage *array*. An expansion cabinet allows up to 18 drive shelves per array. The storage capacity of each disk drive is 36, 72, 146 GB, or higher. The controllers and drive shelves are interconnected in the rack in one of two ways:

- By Fibre Channel loop switches that provide a fault-tolerant physical loop topology.
- Through the use of an expansion panel. If an expansion panel is used, a Fibre Channel loop (controller-to-shelf connection via fiber optic cable) is made by connecting two or four shelves directly to the HSV controllers, and in turn, linking the rest of the shelves together in a daisy chain fashion.




Figure 2: A typical EVA rack configuration (Model 2C12D)

Fibre Channel SAN switches

Fibre Channel SAN switches (see [Figure 3](#)) provide ports to connect all elements of the fabric. Each port can accommodate shortwave, longwave, or very long distance gigabit interface converters (GBICs). GBICs and small form factor pluggable (SFP) GBICs are optical-to-electrical converters inserted into the ports of the Fibre Channel switches to serve as the interface between the fiber optic cables and the switch. GBICs (1-gigabit per second) have SC connectors. SFPs are GBICs that operate at 2-gigabits per second and have LC connectors.




Figure 3: A typical Fibre Channel SAN switch

Three switch product lines are presently supported: B-series, C-series, and M-series. For information on incorporating these switches into your SAN, consult the *HP StorageWorks SAN Design Reference Guide*. For information on individual switch products, browse to:

<http://h18006.www1.hp.com/storage/saninfrastructure.html>.

Storage Management Appliance

The HP OpenView Storage Management Appliance runs software applications at a centralized location for managing and monitoring your storage. The SMA is located out of the data path, which allows data transfers to proceed independently between servers and storage devices. The need for multiple terminal connections for managing and monitoring SAN elements is thus eliminated.

One SMA is required at each site to provide for disaster tolerance. HP recommends the use of two SMAs at each site to also provide high availability. Two Fibre Channel adapters (FCAs) in the SMA provide redundancy. All SMAs can be managing storage systems, but only one SMA can manage a storage system at a time. Up to 16 Continuous Access EVA storage systems can be managed by a single SMA. One SMA must manage the source and destination storage systems when Continuous Access EVA relationships are initially created. Afterwards, these storage system can be managed by separate SMAs.

Fibre Channel adapters

FCAs are inserted into the available slots on a host computer or SMA bus to establish communication between the device and the fabric. A Fibre Channel connection is made by inserting a multimode fiber optic cable from each FCA to an individual port on the Fibre Channel switch.

Required software

This section describes the software applications needed for Continuous Access EVA.

Virtual Controller Software

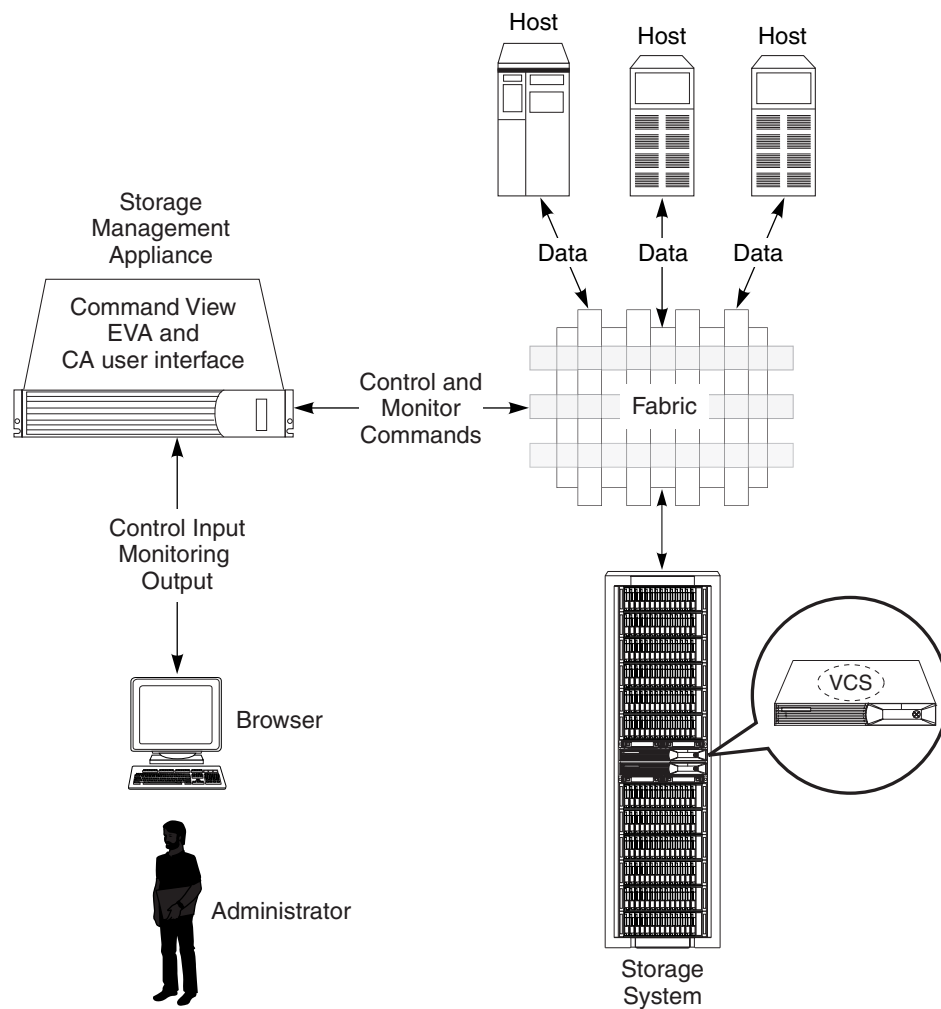
HP StorageWorks Virtual Controller Software (VCS) provides the functionality for the HSV array controller. [Figure 4](#) illustrates the role of VCS in a single storage system. For additional information, refer to the *HP StorageWorks Enterprise Virtual Array User Guide*.

Storage Management Appliance software

The SMA arrives preloaded with SMA software upon which other management applications can be loaded. HP StorageWorks Command View EVA and HP StorageWorks Element Manager for HSG software (preinstalled on the SMA), provide physical and logical views of your storage system through a graphical user interface. While HSV- and HSG-controlled storage may coexist on the same SAN, they cannot be on the same logical fabric when performing data replication, and therefore would have to reside in separate management zones. HSG controllers that are not being used for data replication may reside in the same logical fabric with HSV controllers.

Another specialized application for Continuous Access EVA is called the HP StorageWorks Continuous Access user interface. It is also loaded on the SMA and performs management and monitoring functions like Command View EVA. Its strengths are managing copy sets, data replication groups, and managed sets. The Continuous Access user interface requires the Java Runtime Environment (JRE) application.

These user interfaces are accessed through a Web browser and differ in the way they display various functions and data groups. See Chapter 4 for the recommended uses for Command View EVA and the Continuous Access user interface.



CXO8166A

Figure 4: Functional relationship between controllers and SMA

Host operating systems

Continuous Access EVA supports the following host operating systems:

- HP OpenVMS
- HP Tru64 UNIX
- HP-UX
- IBM AIX
- Microsoft Windows NT, Windows 2000, and Windows 2003
- Novell NetWare
- Red Hat Linux
- Sun Solaris
- SuSE Linux

For the supported versions of these operating systems, refer to the *HP StorageWorks Continuous Access EVA Release Notes*.

Secure Path

Secure Path is a host- or server-based software application that is required for those operating systems (HP-UX, IBM AIX, Linux, Microsoft Windows NT/2000/2003, Novell NetWare, and Sun Solaris) that do not have native multipath support. The HP Tru64 UNIX and HP OpenVMS operating systems have native multipath support and do not use Secure Path software.

Secure Path provides automatic error recovery from connection failures between the host and the storage system through the use of redundant Fibre Channel paths. If any component in the path between the host and storage system fails, Secure Path immediately redirects all pending and subsequent I/O requests from the failed path to the alternate path, preventing an adapter, cable, or controller failure from disrupting data access.

Licensing

Several types of capacity-based licenses are necessary or optional for Continuous Access EVA:

- Business Copy EVA (optional)
- Continuous Access EVA (required)
 - EVA3000 storage systems
 - EVA5000 storage systems

For specific instructions on acquiring these licenses, obtaining license keys, and activating a license, refer to the *HP StorageWorks Continuous Access EVA Getting Started Guide*.

Business Copy EVA license

The Business Copy EVA license provides snapshot and Snapclone functionality on the EVA. This license is optional and ordered for each controller pair by the amount of locally replicated capacity, in terabytes.

Continuous Access EVA license

The Continuous Access EVA license enables remote replication functionality on the EVA. A license is required for each controller pair and separate licenses are needed for EVA3000 and EVA5000 storage systems. The Continuous Access EVA license is issued by utilized capacity, in terabytes, based on the total storage capacity that will be remotely replicated on the EVA.

Concepts

2

This chapter briefly discusses the concepts of controller virtualization and business copy, and introduces important concepts for understanding the workings of Continuous Access EVA. The Enterprise Virtual Array (EVA) storage environment with HSV controller architecture offers enhancements beyond the existing HSG controller architecture. This chapter explains how these improvements help simplify disaster tolerance.

This chapter covers the following topics:

- [Virtualization concepts](#), page 28
 - [Physical vs. virtual storage](#), page 28
 - [Benefits over traditional storage](#), page 28
 - [Virtual RAID types](#), page 29
- [Business Copy concepts](#), page 31
 - [Snapshots](#), page 31
 - [Snapclones](#), page 32
 - [Remote mirrors](#), page 32
- [Continuous Access EVA concepts](#), page 33
 - [Remote data replication](#), page 33
 - [Copy sets](#), page 34
 - [DR groups](#), page 34
 - [Managed sets](#), page 38
 - [Failover](#), page 39
- [Zoning](#), page 41

Virtualization concepts

Virtual storage technology creates a transparent abstraction of storage at the block level. Using block-level mapping techniques, storage virtualization presents host servers with a logical view of storage in the form of virtual disks, while storing the data blocks on physical storage devices in a way that is completely transparent to the servers. Virtualization provides many important new capabilities that simplify storage management and enhance performance. The virtualization concepts discussed below are:

- [Physical vs. virtual storage](#)
- [Benefits over traditional storage](#)
- [Virtual RAID types](#)

Physical vs. virtual storage

The EVA implements virtualization at the storage system level. An EVA storage system consists of a pair of HSV controllers and the array of physical disks it controls. The restrictions of physical disk capacity are eliminated by effective block mapping performed by the HSV controllers. All raw storage is pooled, and virtual disks (or *Vdisks*) that draw their capacity from one of the pools are created. The named group of physical disks from which the Vdisks are created is called a *disk group*. Physical disk devices can belong to only one disk group, and the virtual distribution of storage across the physical disks cannot be seen by the host. Instead, virtual storage capacity is presented to a host as ordinary SCSI logical units (LUNs). Creation and management of Vdisks is done through software on the HP OpenView Storage Management Appliance (SMA) using a Web browser.

Benefits over traditional storage

Virtualization at the storage system level improves performance by up to eight times over the previous generation of controller architecture. The main reason for this improvement is load balancing and distribution across many more spindles. Virtualization allows data to be redistributed across physical disks within a storage pool if an activity occurs that causes a change to the virtual disk data or pool structure. The EVA uses a leveling algorithm to balance performance without interrupting ongoing workloads. This process redistributes each virtual disk's blocks evenly across as many disks as the disk's redundancy type will allow. Because of storage pooling virtualization, the EVA is able to support multiple virtual disks of varying capacity and RAID types within a single storage pool.

The ability to expand virtual disk capacity dynamically, and without application downtime, greatly improves capacity efficiency. If the host supports this, then virtual disk size can be increased without disrupting the application. Administrators can monitor capacity and dynamically allocate capacity from the storage pool in 1-GB increments when needed. However, support for increased capacity size may be limited by your operating system.

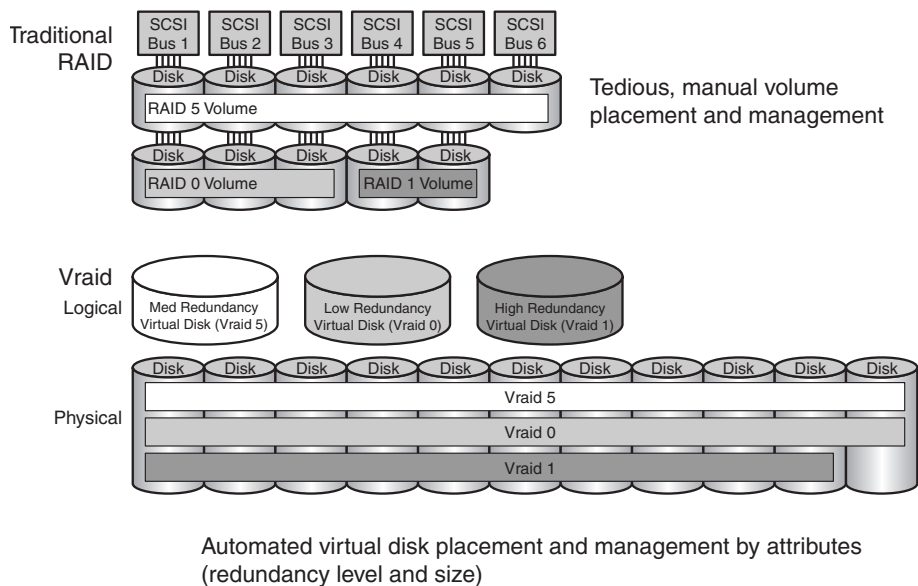
Virtualization simplifies storage management by allowing dynamic attribute changes. Similarly, physical storage components can be added, removed, and managed without the need to manage the host's view of virtual storage and without taking any data offline.

Virtual RAID types

The type of HP VersaStor-enabled virtual RAID (Vraid) used helps optimize your storage system for performance and protection level. Virtual RAID differs from traditional RAID in that the data and its redundancy protection is spread across all of the physical disks in the pool (see [Figure 5](#)). There are three types of redundancy available with the EVA storage system:

- **Vraid0**, or striping, is optimized for speed and virtual disk size but has no parity protection or redundancy. It consumes only one block of physical capacity for every block of usable capacity. Of the three types of Vraid, reading and writing to a striped disk is fastest and makes the fullest use of available storage. Vraid0 is not recommended with Continuous Access EVA because there is no redundancy.
- **Vraid1**, or mirroring, gives the highest level of data protection of the three types of Vraid, but uses the most storage space. Mirrored data is stored twice and thus consumes two blocks of physical capacity for every block of usable capacity. Vraid1 has no restrictions and provides the highest availability.
- **Vraid5**, or parity striping, offers a balance of speed, size, and redundancy when compared to the other two types of Vraid. It consumes one block of parity for every four blocks of data and thus consumes 1.25 blocks of physical capacity for every block of usable capacity. It is less efficient in read performance if a disk failure occurs because lost data must be calculated from the parity and remaining data on other disks. Vraid5 is not recommended for storage systems with less than eight drive shelves, as there is no guarantee of survivability with the loss of a shelf.

HP recommends the use of Vraid1 or Vraid5 with Continuous Access EVA for improved data protection. HP further recommends that the same type of Vraid level be used within a DR group.



CXO7988A

Figure 5: Traditional RAID vs. virtual RAID

The algorithms that control Vraid1 and Vraid5 data layout work best if all the disks in the disk group have the same physical capacity and performance. Disk group efficiency also improves if the group has an even number of disks and is best if the number of disks in the group is an even multiple of eight. At the time of creation, the minimum number of drives in a disk group is eight.

Business Copy concepts

The term *business copy* refers to protected-copy volumes that are created and maintained without interrupting access to the source volumes. HP StorageWorks Business Copy EVA provides the capability to create instantaneous copies of data for development, testing, or backup, without taking your system offline. Three business copy methods are discussed below.

- [Snapshots](#)
- [Snapclones](#)
- [Remote mirrors](#)

Snapshots and Snapclones require a Business Copy license. Remote mirrors require a Continuous Access EVA license.

Snapshots

A *snapshot* is a temporary Vdisk that reflects the contents of another Vdisk at the particular point in time that the snapshot was created. The created Vdisk is always linked to the original Vdisk, and may be created at either the primary or the remote site. Cross-Vraid functionality allows you to make a snapshot and specify a different Vraid type than the original.

Two types of snapshots are supported: fully allocated (the default) and demand allocated.

With a fully allocated snapshot, a set amount of capacity equal to the original volume is reserved. Data is not written to that reserved space until necessary. As the data changes in the original virtual disk, the data in the snapshot volume is updated to mirror the original volume. You may desire to make a fully allocated snapshot when a significant portion of the data will be changing over time, or when the snapshot itself will remain on the storage system for extended periods of time.

In contrast, the demand-allocated snapshot does not reserve capacity for the snapshot volume in advance. Rather, space on the snapshot volume is used only as the original virtual disk's data changes. The snapshot volume is a new virtual disk that initially shares the original virtual disk's pointer-based entries. As new data is written to the original virtual disk, the old data is copied to preserve the original contents of the snapshot. The allocation on demand snapshot is especially useful when only a small portion of the virtual disk is expected to change over time, or in situations where the snapshot will exist for only a short period of time before a backup procedure occurs. A significant feature of both types of snapshots is that they can be created from any level of redundancy (Vraid0, Vraid1, or Vraid5).

Snapclones

A *Snapclone* is a logical disk that is an exact copy of another logical disk at the particular point in time that it is created. The link from the Snapclone to the original Vdisk is dissolved upon the completion of the full data copy. A Snapclone may be created at either the primary or the remote site, and a different Vraid type can be specified upon creation.

With the virtually instantaneous Snapclone, a complete copy of the original virtual disk is made as quickly as data transfer rates permit. Therefore, the Snapclone is the best option for preserving a long-term copy or a series of copies of a virtual disk. Using snapshot technology, all data is copied into the reserved space proactively so the result is two identical copies of the data, at the redundancy level of the original volume, in the shortest time possible.

This process is unlike traditional cloning methods where the clone copy is not available until the copy is complete. As the Snapclone is created, the controller is able to access the original virtual disk for the data and keep track of what data has changed since the moment the Snapclone was taken. The benefit of a Snapclone is that you get an immediate point-in-time copy of a virtual disk that can be presented to a server.

Remote mirrors

Continuous Access EVA allows Vdisks to be remotely mirrored across a long-distance link. The mirroring options or link type used (for example, Fibre Channel over Internet Protocol, wavelength division multiplexing (WDM), and so on) can be any of those currently supported by Continuous Access EVA. Consider remote mirroring for the following purposes:

- **Disaster recovery**—Maintaining a “failover” datacenter.
- **Data migration**—Moving data from one storage system or data center to another.
- **Data distribution**—Pushing copies of data between geographically dispersed storage systems.

Continuous Access EVA concepts

Continuous Access EVA uses the remote-copy function of the HSV controller running VCS V3 or later to achieve host-independent data replication. The remote copy is the major feature of the Continuous Access EVA solution. Storage system management is provided by applications residing on the SMA. This section describes some basic Continuous Access EVA terminology, concepts, and features. An understanding of these topics is necessary before you begin planning and implementing your solution. The topics discussed are:

- [Remote data replication](#)
- [Copy sets](#)
- [DR groups](#)
- [Managed sets](#)
- [Failover](#)

Remote data replication

The HSV storage system at the primary location is connected to a partner storage system at the alternate location. To create data replication for storage, a source Vdisk is configured at the primary storage system. When data replication is selected, the destination Vdisk is automatically created by software at the remote storage system. Any data written to the source Vdisk is then mirrored to the destination Vdisk. Applications continue to run while data replication goes on in the background over a separate interconnect. When a storage system contains both source Vdisks and destination Vdisks, it is said to be *bidirectional*. A storage system can have a bidirectional data replication relationship with up to two other storage systems, and an individual Vdisk can have a unidirectional replicating relationship with only one other Vdisk.

The remote copy feature is intended not only for disaster recovery but also to replicate data from one storage system or physical site to another storage system or site. It also provides a method for performing a backup at either the source or destination storage systems.

Copy sets

Vdisks are user-defined storage allotments of virtual or logical data storage. A pairing relationship can be created to automatically replicate a logical disk to another logical disk. The generic term for this replication is a *copy set*. A *relationship* refers to the arrangement created when two storage systems are partnered for the purpose of replicating data between them. Copy sets are part of a larger construct called a DR group, which is described in the next section.

A Vdisk does not have to be part of a copy set. Vdisks at any site can be set up for local storage and used for activities such as testing and backup. Clones and Snapclones are examples of Vdisks used in this manner. When a Vdisk is not part of a copy set, it is not disaster tolerant, but it can use various Vraid types for failure tolerance.

DR groups


A *data replication (DR) group* is a VCS construct comprising one or more Vdisks in an EVA storage system so that they:

- Replicate to the same specified destination storage system
- Fail over together
- Preserve write order within the collection
- Share a log disk

All virtual disks used for replication must belong to a DR group, and a DR group must contain at least one Vdisk on separate storage systems. A DR group can be thought of as a collection of one to eight copy sets, however for optimal performance, each DR group should be limited to one Vdisk. [Figure 6](#) depicts the replication of one DR group between separate sites.

The replicating direction of a DR group is always from a source to a destination. By default, the storage system on which the source Vdisk is created is called the *Home* storage system. The Home designation denotes the preferred storage system for the source and this designation can be changed to another storage system. The concept of Home only exists within the context of the Continuous Access user interface and not for Command View EVA.

A DR group replicating from a Home storage system to a destination system is in the *original* state. When replication occurs from a storage system that was created as the destination to the Home storage system (for example, after a failover, which is discussed later), it is in a *reversed* state.



CXO7989A

Figure 6: Continuous Access EVA DR group replication

All members of a DR group are presented to a host through the same controller and move together. Therefore, you want the same preferred path selected for all DR group members, with presentation to same Fibre Channel adapters (FCAs). With clustered hosts, this DR group presentation is particularly important.

DR group properties

Properties are defined for every DR group that is created. Some DR group properties are described below:

- **Name**—A unique name given to each DR group. HP recommends that the names of replicating DR groups at the source and destination be the same.
- **DR Mode**
 - **Source**—A DR group established as an active source that replicates to a passive destination.
 - **Destination**—A DR group established as a passive destination that receives replication data from an active source.
- **Failsafe mode**—When this mode is enabled, all source Vdisks within the DR group become both unreadable and unwritable if any member becomes unreachable. This condition is known as *failsafe-locked* and requires immediate intervention. When the failsafe mode is disabled and the destination Vdisk is unreachable, normal logging occurs.
- **Connected system**—A pointer to the storage system where the DR group is replicated.
- **Write mode**
 - **Asynchronous mode**—A write operation provides an I/O completion acknowledgement to the host after data is delivered to cache at the source controller, but before data delivery to cache on the destination controller.
 - **Synchronous mode**—An I/O completion acknowledgement is sent to the host after data is written to the source and destination caches.
- **Suspension**
 - **Suspend**—When this command is issued and failsafe mode is not enabled, I/O replication is halted between the source and destination Vdisks. Source Vdisks continue to run I/O locally and the I/O is also copied to the DR group log Vdisk. The suspend command is not available if failsafe is enabled.

- **Resume**—When this command is issued, replication resumes between the source and destination Vdisks. Merging of the log Vdisk or a full copy is also performed.

Log disks

The DR group *log* is Vraid1 storage that is allocated on demand to collect host write commands if access to the destination storage system is severed. When a connection is later re-established, the contents of the log are written to the destination Vdisk to synchronize it with the source Vdisk. This process of writing the log contents, in the order that the writes occurred, is called *merging*. Sometimes it is more practical to copy the source Vdisk directly to the destination Vdisk. This copy operation is called a *full copy*—all 1-MB blocks written on a source Vdisk since it was created are copied to the destination Vdisk. There is no manual method for forcing a full copy, rather it is an automatic process that occurs when a log is full.

A log can be in one of three states:

- **Normal**—No source Vdisk is logging or merging.
- **Logging**—At least one source Vdisk of the DR group is logging (capturing host write commands) but none are merging.
- **Merging**—At least one source Vdisk of the DR group is merging and logging.


When a DR group is in a logging state, the log will grow in proportion to the amount of write I/O being sent to the source Vdisks. As the log grows, more space is allocated out of the available capacity of the disk group where it is a member. The capacity available to the log does not include the spare capacity or any capacity being used by Vdisks or Snapclones. This means the log disk will never overwrite any other data. Similarly, when a DR group is logging, the available capacity for creating Vdisks, snapshots, and Snapclones does not include the capacity already used by the log disk. Therefore, a log disk will never be overwritten by any other data.

When creating disk groups and distributing Vdisks within them, sufficient capacity must remain for log disks to expand to their maximum level. A log will be placed into any near-online disk group if one exists; otherwise, it will be placed into an online disk group with the most average free space. The log is declared full, and reaches its maximum level, whenever the first of the following conditions is reached:

- The size of the log data file exceeds twice the capacity of the DR group.
- No free space remains in the physical disk group.
- The log reaches 2 TB of Vraid1 (4 TB total).

Managed sets

A *managed set* is a collection of DR groups selected by the user for the purpose of performing the same operation on them. For example, a managed set can be created to manage all DR groups of a particular set of applications that reside in separate storage systems. Managed sets are only available through the Continuous Access user interface. [Figure 7](#) depicts a managed set that contains common DR groups on a source storage system that replicate to a destination storage system.



CXO7990A

Figure 7: Managed sets

DR groups can belong to more than one managed set and managed sets may include DR groups from more than one storage system. Managed sets are only useful if managed by the same SMA at the time a failover is performed.

Managed sets are a useful way to reduce the number of steps needed to perform supported operations on the members in the set. It does not provide support for “atomic” operations (defined as the stopping of I/O to all members of the managed set, performing selected operations, and the restarting of I/O). Instead, a command is sent to each member in the order they were added to the managed set, but the command does not proceed to the next member until it completes.

Failover

The recovery process whereby one DR group, managed set, fabric, or controller switches over to its backup is called *failover*. The process can be planned or unplanned. A planned failover allows an orderly shutdown of the system before the redundant system takes over. An unplanned failover occurs when a failure or outage occurs that may not allow an orderly transition of roles.

Listed below are several types of Continuous Access EVA failovers:

- **DR group failover**—An operation to reverse the replication direction of a DR group. A storage system can have a replication relationship to two other storage systems, but a DR group can only replicate between two storage systems. [Figure 8](#) shows a data replication relationship among DR groups at three locations. Storage systems A and D are located at the primary site and storage systems B and C are located at two alternate sites. If contact with the primary site is broken, failover can manually occur with the destination DR groups that were replicating to storage system B. Storage system B then acts as the primary site for these DR groups after failover. At storage system C, the source DR groups begin logging until new remote sites are re-established.
- **Managed set failover**—An operation to reverse the replication direction of all the DR groups in the managed set.
- **Fabric or path failover**—The act of transferring I/O operations from one fabric or path to another.
- **Controller failover**—The assumption by one controller of the workload of its partner (within the same storage system).





Figure 8: Replicating relationships among DR groups

Zoning

Zoning is a logical grouping of end-to-end Fibre Channel connections, implemented by switches, to create a barrier between different environments and allow for finer segmentation of the fabric. Switch ports that are members of a zone can communicate with each other but are isolated from ports in other zones.

Because the certain hosts in a Continuous Access EVA environment can conflict with each other, they must reside in separate zones. If any HSG80 controllers running Data Replication Manager (DRM) reside within the SAN, they must be reside out of the management zone of the SMA. HSG80 controllers that are not running DRM can reside inside the same SMA management zone as the EVA.

[Figure 9](#) shows an example of zoning in a simple configuration that keeps the SMA and hosts separated but allows each access to EVA storage.



CXO7994A

Figure 9: Simple Continuous Access EVA zoning

Configuration Planning

3

This chapter provides suggestions and strategies for planning a Continuous Access EVA configuration. Initial planning allows you to get the most benefit from a well-designed SAN. Planning can reduce recovery times and associated downtime costs due to unexpected outages, especially by reducing the confusion that can occur during disaster recovery. Having a documented design plan in place that shows names, identifiers, and connections used in your configuration expedites the planning, setup, failover, and recovery process. Your plan should be up-to-date with copies at all sites.

The topics discussed in this chapter specifically address configuration planning, which is done before you begin to create replicating virtual disks for your data. For higher level planning recommendations, refer to the *HP StorageWorks Continuous Access EVA Design Reference Guide*.

This chapter covers the following topics:

- [About disk groups](#), page 44
- [Planning your disk group configuration](#), page 46
- [Naming restrictions](#), page 48
- [Planning fabric configurations](#), page 49
 - [Dual-fabric configuration](#), page 49
 - [Single-fabric configurations](#), page 50
- [Load balancing](#), page 52
- [Planning your zones](#), page 53
- [Restrictions](#), page 56

About disk groups

Initialization is a process required to enable the storage system to be used. Initialization binds the controllers together as an operational pair and establishes preliminary data structures on the disk array. Initialization also sets up the first disk group, called the default disk group.

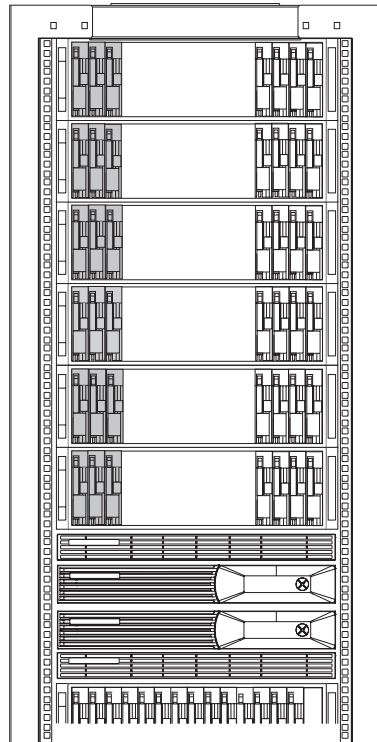
The physical disks connected to each controller pair must be assigned to a *disk group*: a named grouping of at least eight physical disks from which virtual storage is created. Each disk group constitutes a separate storage pool and is independent of other disk groups in the storage system.


Before you begin the initialization process, perform an analysis of what types of data will reside on your storage system. Determine what purpose your drives will serve (for example, performance vs. capacity) so they can be allocated to separate disk groups. For example, a 36-GB, 15,000-rpm drive is often used for performance, while a 300-GB, 10,000-rpm drive is better suited for capacity. Fibre Attached Technology Adapted (FATA) drives provide larger capacities at a lower costs, but with slightly lower random access times. A disk group must contain only one disk drive model or you risk a performance impact. VCS will restrict disk group membership to only online or near-online (FATA) class drives.


When you create a disk group, you cannot select which disks will be part of the group unless you build the disk groups sequentially. Begin initializing your storage system by adding only those physical drives required for your first disk group, and then create the disk group. Distributing the disks among the shelves reduces the exposure to failures in the disk backplane. Next, add more disk drives to the storage system for a second disk group, and then create that disk group. [Figure 10](#) shows the sequential building of disk groups in this vertical fashion (among the shelves).

As part of the disk group creation process, you must select a disk failure protection level. Three choices are possible: None, Single, and Double. The disk drive failure protection level represents the amount of reserved capacity allocated on the storage system for data recovery in the event of a failed or failing disk drive. The reserved capacity is based on the largest drive in the disk group. The system must cover a failure in any drive, so it reserves enough capacity to cover the largest failure that could happen. HP recommends the Double option so as to provide the most data protection. Selecting the Double option requires the largest amount of disk drive capacity set aside to provide this protection.

For example, if a member of a Vraid1 mirror pair fails, both members are marked as failed, and the capacity of both are replaced. So in this case, a single protection level means that two times the capacity of the lost member is replaced. A double protection level means that four times the capacity of the largest Vraid1 drive was reserved for a failure.



 = Drives installed with
1st group

 = Drives installed with
2nd group

CXO8167A

Figure 10: Building sequential disk groups

Planning your disk group configuration

The recommended steps for planning the configuration of a disk group are as follows:

- a. If performing bidirectional replication, the disk groups must be symmetric on both storage systems.
1. Evaluate the performance requirements and the I/O pattern.
 - a. If the storage system I/O stream is dominated by I/Os of substantially equivalent transfer size and locality (for example, mostly sequential big transfers or mostly random small transfers), or does not require very high sustained sequential throughput, continue to step 3.
 - b. If the application I/O stream is dominated by a mix of simultaneous sequential and random transfers, determine how these streams need to be directed to specific Vdisks. Must they all be directed to the same Vdisk, or can they be separated to different Vdisks?
 - c. In general, separate sequential I/O stream data (database logs, rich content) from random I/O streams (database information store, file shares). That is, put like-task virtual disks into the same disk group (I/O stream types).

Note that transfer profiles to a given disk that differ over time are not a major consideration. A Vdisk that receives sequential transfers for part of the day and random accesses for the rest of the day operates well in both cases. The situation of interest occurs where simultaneous sequential and random streams must be accommodated.

2. Determine the minimum number of disk groups needed. Remember that having too many disk groups may strand spare capacity. Stranded capacity is physical capacity that cannot be utilized because it cannot meet the structural or other requirements defined for a new Vdisk. Conversely, a disk group with a large number of members will perform better than two smaller groups, each with half the number of members. DR groups may contain Vdisks from multiple disk groups, but must be on the same storage system.
 - a. Create the fewest number of disk groups consistent with the failure and performance isolation requirements. One disk group that consists of all the physical drives may be the best choice for most users.
 - b. If the storage system requires the “vertical” disk groups described in step 1, the number of groups is already determined, and the configuration choice is limited to 2, 6, 8, 12, or 18 shelves, depending on capacity needs.

Note that in this configuration, it is acceptable to have fewer than 14 groups; this is accomplished by reducing the number of drives on each shelf. All shelves should have approximately the same number of drives.

- c. If the storage system does not require “vertical” disk groups, the number of groups should be the greater of the number of groups required to get reliability separation or the number of groups required for performance locality. For example, if the storage system needs four independent failure domains, but requires only that one sequential stream be separated from all the others, then configure four groups. Conversely, if the installation requires only two failure domains, but has six sequential I/O streams, configure for six groups.
3. Determine the number of disks in each group.
 - a. Always configure groups with drives of identical or similar capacity and performance.
 - b. Be generous with the number of disks in each group, but keep it an even multiple of 2 and strive for the best availability of disk groups with drive counts that are multiples of 8—that is, 8, 16, 24, 32, and so on. The minimum is 8 drives per group and the maximum is 240 drives.
 - c. After identifying the number of groups, determine which virtual disks are to reside in each group, and then size each group for the appropriate capacity, given the above guidelines. Keep in mind that each group needs its own spare capacity. Also, factor in the snapshot, Snapclone, and write history log needs. A formula for determining the number of drives for a given capacity is in the *HP StorageWorks Continuous Access EVA Design Reference Guide*.
 - d. When sizing groups, also consider that the efficiency of the controller virtualization algorithms improves with the amount of free space available. Given the trade-off between the cost of disk capacity and the cost of human management time, extra capacity almost always pays off in lower management overhead and typically higher peak performance. A good planning value is to estimate about 10 percent for free space.

Note: For the best long-term performance, background controller processes require at least 10 percent free space.

Naming restrictions

Two categories of names can be created for the EVA using a management user interface:

- **Storage system names**—A user may enter 20 characters to name a storage system. This is the number of characters that can appear on the operator control panel. Some special characters are allowed.
- **Normal object names**—These are names given to Vdisks, disk groups, DR groups, folder names, and so on. The limit is 32 characters, and many printable special characters are allowed. Consecutive spaces are not allowed.

Normal object names are case-sensitive. For example, you may have two Vdisks named Disk1 and disk1, although it is not recommended.

There are restrictions to some of the characters used for both name categories. The storage system names can contain any printable character except for the following:

? (question mark)	" (double quotes)	/ (slash)
\ (backslash)	< (less than sign)	> (greater than sign)
* (asterisk)	(vertical bar)	: (colon)
% (per cent)	& (ampersand)	, (comma)
+ (plus)	Consecutive spaces	

The comment fields that reside in the object properties can use all the printable characters. All comment fields allow up to 128 characters, with the exception of the comment field for DR groups. The DR group comment field is limited to 64 characters.


Folders provide a convenient way to group and organize data when using the Command View EVA or the Continuous Access user interface. Placing Vdisks, hosts, replication groups, and storage systems into folders identified by easily understandable and recognizable names and descriptions can help you and others quickly find and manage data when there is an urgent need. Ensure that everyone that may be involved with disaster planning is aware of chosen folder names and the type of information in each folder.

Planning fabric configurations

Continuous Access EVA uses controller replication to provide disaster-tolerant data across a fabric provided by Fibre Channel switches. Switches can be arranged in different topologies, depending on your needs or limitations, such as distance, simplicity, or cost. For a discussion of various topologies, refer to the *HP StorageWorks SAN Design Guide*. Although a dual-fabric configuration with at least four switches and two Fibre Channel adapters (FCAs) per host is the minimum recommended for high availability, alternate configurations are available for users who are not concerned about interruptions due to points of failure in the fabric.

Dual-fabric configuration

To protect data and provide high availability, two fabrics using multiple switches are necessary, as shown in [Figure 11](#). Each fabric provides a route through an interswitch link (ISL) to connect your sites. If one fabric becomes inoperable because of a hardware or software problem, the other fabric can assume the additional traffic load until the problem is resolved. A Storage Management Appliance (SMA) is required at each site for redundancy and high availability. If data protection and minimal downtime is of primary importance to your data center, then a dual-fabric configuration is necessary.



CX08218A

Figure 11: Dual-fabric configuration

A dual-fabric configuration may also allow you to increase the maximum allowable port count. If your configuration requirements exceed the allowable port count of a single fabric, determined by the type of switches, a second fabric enables the total port count to be increased.

Two FCAs per host allows a path to both fabrics. A reduced configuration with one FCA in a host, connected to one fabric, still allows replication to occur. However, this introduces a single point of failure, and does not provide a highly available configuration. Run multipathing software (Secure Path) on those hosts where multipathing is not built into the operating system, so that fabric redundancy is available.

Single-fabric configurations

A single-fabric configuration is intended as an entry-level, or proof-of-concept solution. It may address your needs in a small environment by permitting testing or data transfers to occur with minimal equipment and cost. The configuration is easy to manage because only one zoning configuration needs to be maintained. However, any mismanagement in a single fabric can cause fabric downtime. Redundant SMAs are not necessary in single fabrics because of the lower expectations for high availability.

The simplest single-fabric configuration is one that uses one switch as shown in [Figure 12](#). This single-switch fabric maximizes performance because every switch port has full connectivity to every other port on the switch. This configuration can be used to test most Continuous Access EVA features, but not those that depend on fabric redundancy. Because no fabric redundancy is provided, this solution is neither highly available nor disaster tolerant.





Figure 12: A single-switch fabric configuration

A dual-fabric configuration can be created for testing purposes using a single switch that has been zoned to create two logical fabrics. However, this configuration is only as reliable as the single switch being used, and is therefore not considered highly available.


Multiple switches can be used to create single-fabric configurations. [Figure 13](#) shows a configuration with more flexibility than one using only a single switch. With switch separation provided by an ISL between two switches, longer distances can be achieved, and testing can include the loss of a link or a site. Two SMAs are recommended with site separations and multiple SMAs can be tested with this configuration. Given an appropriate amount of site separation, this solution may be disaster tolerant.



CXO8220A

Figure 13: Two-switch, single-fabric configuration

Another multiple-switch, single-fabric configuration makes use of a mesh switch fabric, like that shown in [Figure 14](#). An advantage of this configuration over a dual-redundant fabric is a higher resiliency. The mesh design allows the fabric to heal itself in the event of a switch, port, or path failure, and the EVA and hosts experience minimal interruption. Static routing or link cost associations may be used to optimize load balancing across the ISLs. This configuration provides disaster tolerance if the four site-to-site links are long enough, and a reduced level of high availability as perturbations can ripple across a single fabric. Again, only dual-redundant fabrics provide the highest levels of availability.



CXO8219A

Figure 14: Mesh-switch fabric configuration

Hosts with a single FCA can be run in any of these single-fabric environments provided that a multipath driver is also used (Secure Path software or native multipathing). Data replication can occur, but in addition to reduced availability introduced by the single fabric, less availability and another point of failure is introduced by the single host FCA. These limitations can be mitigated with server clustering or n+1 sparing, but it still allows multiple single points of failure in the configuration.

Load balancing

Continuous Access EVA works best when the average workload (reads and writes) is applied somewhat equally to both controllers, and therefore to both fabrics and intersite links. To obtain this balance, manual measurements should be made in an attempt to keep the utilization rate of either intersite link below 40 percent, so that if one link fails, the average utilization of the surviving link does not exceed 80 percent. Similarly, the utilization rate of a single controller as measured on both host ports should not exceed 45 percent on average, or peak above 50 percent, to prevent overloading a surviving controller should one of them fail.

There are two ways to balance the workload:

- Let the hosts make the arrangements
- Prior planning when setting up the replicating LUNS

Letting the hosts balance the workload is acceptable if there are only a few hosts, and they all support some type of auto-load balancing with Secure Path. Beyond this scope, and because the hosts do not share workload information with other hosts, it is best to plan ahead and use the default load balancing tools of the EVA. Specifically, all members of a DR group must be preferred to the same controller and FCA port pair because they must perform the same actions together as a group. For additional information see “Creating Vdisks” on page 70.

Planning your zones


All HP Fibre Channel switches provide a fabric zoning feature. The use of zoning is required under the following specific conditions:

- Use zoning when combining different hardware platforms, operating systems, or storage systems that are currently supported only in homogeneous SANs, and it is unknown whether there are interaction problems. [Table 3](#) shows the zone compatibility for different platforms in a Continuous Access EVA SAN. Platforms in the same column can exist in the same SAN.

Table 3: Continuous Access EVA Platform Zoning Requirements

Zone 1	Zone 2	Zone 3	Zone 4	Zone 5
HP OpenVMS	HP OpenVMS	Linux	HP-UX	IBM AIX
HP Tru64 UNIX	HP Tru64 UNIX			
Microsoft Windows NT/2000/2003	Microsoft Windows NT/2000/2003			
Novell NetWare	Sun Solaris			

- Use zoning when there are known interaction problems between different hardware platforms or operating systems and specific storage types. For example, the SMA must not share the same zone with any hosts or HSG80 controllers running Data Replication Manager (DRM). For a listing of platform interoperability issues by operating system version number, consult Chapter 4 of the *HP StorageWorks SAN Design Reference Guide*.
- Use zoning when the number of nodes or ports in the SAN exceed support limits for storage system connections. There is a connection limit for storage systems using the HSG or HSV controllers. The version of ACS or VCS controller code determines the specific limit.
- Use zoning to limit the management scope of SMAs. Continuous Access EVA limits a single management zone to 16 arrays. See the *HP StorageWorks Continuous Access EVA Design Reference Guide* for management versus distance considerations.
- Use zoning to isolate an SMA from another SMA. [Figure 15](#) shows a zoning configuration where SMA A can manage the storage at the primary and alternate sites. SMA B is also zoned to manage both sites in the event of failure at the primary site. Either SMA can be used to manage a storage system, but only one can manage a storage system at a time.



CXO8173B

Figure 15: SMA management zones

HP suggests that you use [Table 4](#), the zoning input form, to capture and track the required device and command information. The form supports two paths, two switches, and a maximum of 16 entries per switch. Copy and use a separate form to track information for each zone. Organize alias name, function, and site data by either World Wide Name (WWN) number or port ID number.

Use the host WWN when zoning the storage system controllers. Each host controller port connects to a separate fabric, and sometimes the designations can switch between the controller pairs.

Table 4: Zoning input form**Zoning Configuration Name =**

Zone Name=

Switch Name=

Path=

WWN #	Domain ID #	Port #	Alias Name	Function	Site

Zoning Configuration Name =

Zone Name=

Switch Name=

Path=

WWN #	Domain ID #	Port #	Alias Name	Function	Site

Restrictions

Table 5 lists restrictions to consider before configuring your Continuous Access EVA solution.

Table 5: Restrictions

Restriction	Comments
Maximum of 16 storage systems on the same SAN.	
A disk group must contain at least eight physical drives.	Disk drives should be even multiples of the number of drive enclosures and 8.
Each storage system can support 240 physical disk drives.	
One DR group supports a maximum of 8 copy sets from one or more disk groups within the same storage system.	
A storage system can support up to 512 Vdisks in a storage pool.	This includes copy sets, non-copy Vdisks, snapshots, and Snapclones.
One storage system supports a maximum of 128 one-member DR groups.	Multiple member DR groups reduce the number of DR groups allowed. For example, if each DR group contained 8 Vdisks, then at most 16 DR groups can be created.
Vdisk capacity can range from 1 GB to 2 TB.	
A total of 256 FCAs are allowed per storage system.	Port count consists of single and dual-port FCAs and does not have to be symmetric between storage systems.
Each storage system cannot contain more than 32 TB of presented virtual disks capacity.	
Each storage system can have a relationship with one or two other storage systems.	
HSG controllers may be present on the SAN but cannot interoperate with HSV controllers.	The HSG and HSV controllers must be in different management zones. HSG80 controllers running DRM cannot be in any SMA management environment.
Up to 7 snapshots or Snapclones are allowed per Vdisk.	Can be on either the source or destination

Table 5: Restrictions (Continued)

Restriction	Comments
Up to 8 snapshots or Snapclones are allowed per DR group.	
A site-to-site latency of 100 ms or less	Refer to the <i>HP StorageWorks Continuous Access EVA Design Reference Guide</i> for distance limitations.
Each Continuous Access EVA management zone must have one active and one standby SMA.	
Maximum of 7 hops per Fibre Channel fabric between controllers for B-series switches.	
Maximum of 3 hops per Fibre Channel fabric between controllers for C-series switches.	Maximum of 2 hops in a non-Continuous Access EVA environment.
Maximum of 3 hops per Fibre Channel fabric between controllers for M-series switches.	
Maximum of 28 B-series, 11 C-series, or 24 M-series switches per fabric.	Refer to the <i>HP StorageWorks SAN Design Reference Guide</i> for interoperability rules between B-, C-, and M-series switches.
Qualified operating systems are HP Tru64 UNIX, HP OpenVMS, HP-UX, IBM AIX, Linux (Red Hat and SuSE), Microsoft Windows NT/2000/2003, Novell NetWare, and Sun Solaris	Refer to the <i>HP StorageWorks SAN Design Reference Guide</i> for specific zoning rules.

Configuration

4

This chapter discusses how to configure Continuous Access EVA hardware and software. Because Continuous Access EVA spans multiple sites, you must configure the systems at each site.

This chapter covers the following topics:

- [Hardware configuration](#), page 60
 - [Fibre Channel adapter installation](#), page 60
 - [Configuring the Fibre Channel switches](#), page 61
 - [Controller-to-switch connections](#), page 61
 - [Host-to-switch connections](#), page 62
 - [EVA zoning recommendations](#), page 62
 - [HSG80 zoning recommendations](#), page 64
 - [SMA-to-switch connections](#), page 65
- [Software configuration](#), page 65
 - [Storage Management Appliance software setup](#), page 66
 - [Configuring hosts](#), page 68
 - [Licensing](#), page 69
 - [Initializing your storage systems](#), page 69
 - [Creating disk groups](#), page 69
 - [Creating host folders](#), page 70
 - [Creating hosts](#), page 70
 - [Creating Vdisk folders](#), page 70
 - [Creating Vdisks](#), page 70
 - [Installing Secure Path](#), page 72
 - [Presenting Vdisks to hosts](#), page 72

- [Accessing the Continuous Access user interface](#), page 72
- [Creating copy sets and DR groups](#), page 75
- [Presenting a copy set to a destination host](#), page 79
- [Specifying disk group membership for a log](#), page 79
- [Deleting or detaching copy sets](#), page 81
- [Deleting DR groups](#), page 81
- [Creating managed sets](#), page 81
- [Editing a managed set](#), page 82
- [Adding a DR group to a managed set](#), page 82
- [Removing a DR group from a managed set](#), page 83
- [Deleting a managed set](#), page 83
- [Backing up configuration information](#), page 84

Hardware configuration

The following sections describe hardware configuration tasks that must occur before you begin the software configuration.

Fibre Channel adapter installation

Fibre Channel adapters (FCAs) must be installed in each host. For high availability, at least one dual-port or two single-port FCAs are needed, with up to three pairs of any combination of single- or dual-port FCAs allowed per host. The section “Planning fabric configurations” on page 49 describes several instances where a single FCA can be used when high availability is not expected. For detailed FCA installation instructions, refer to the documentation that comes in the host kit or with your adapter.

Locate and record the World Wide Names (WWNs) of each Fibre Channel adapter on the zoning worksheet provided in Chapter 3. Keep a copy of the worksheets at all your sites. In addition, you will need to record the WWNs for the storage arrays and the Storage Management Appliances (SMAs) for each site.

Note: The WWN can be found on the bottom of the Fibre Channel adapter board. Look for a small bar code label with an IEEE (Institute of Electrical and Electronics Engineers) precursor. A WWN example is 1000-0000-C920-A5BA. The WWN for an EVA is located on the front of the rack at the sides of the controllers.

Configuring the Fibre Channel switches

Your Fibre Channel switches must be installed and configured with two working redundant fabrics before you connect the remaining Continuous EVA components to your fabrics. The section “[Planning fabric configurations](#)” on page 49 describes several instances where single fabrics can be used when high availability is not expected. For information on the specific switches used and gigabit interface converters (GBICs) needed, refer to the following website:

<http://h18006.www1.hp.com/storage/saninfrastructure.html>

Controller-to-switch connections

Before connecting fiber optic cables between storage components, HP recommends that you tag each end to identify switch names, port numbers, controller names, and so on.

Four fiber optic cable connections are required for each storage array. The only supported connection scheme is shown in [Figure 16](#). Connect the fiber optic cable such that port 1 of Controller A and Controller B go to different fabrics. Connect port 2 of Controller A and Controller B to separate fabrics that are the fabric opposite from port 1 on that controller.

Either controller can be Controller A or Controller B. In a storage system that has not been configured, the first controller that powers up and passes self-test becomes Controller A. This controller designation is remembered and remains upon subsequent power ups.




Figure 16: Controller-to-switch cabling

Host-to-switch connections

Tag each end of your fiber optic cable to identify switch names, port numbers, host names, and so on.

Two fiber optic connections are required for each pair of FCA ports in the host. Connect the fiber optic cable such that each cable in the pair of FCAs goes to separate fabrics.

EVA zoning recommendations

The following EVA zoning recommendations are provided for the three switch product lines: B-series, C-series, and M-series.

Zoning with B-series switches

With B-series switches, the following naming conventions apply:

- The storage system WWN ends with a 0.
- The Controller A, port 1 is the storage system WWN but ending with a 9.
- The Controller A, port 2 is the storage system WWN but ending with an 8.

- The Controller B, port 1 is the storage system WWN but ending with a D.
- The Controller B, port 2 is the storage system WWN but ending with a C.

HP recommends that you zone using the host WWN address for each fabric instead of the controller host port WWN.

For example, in [Figure 17](#), the storage system host WWN is designated as 50:00:1f:e1:00:15:40:80. Cabled to this fabric are Controller A, port 2 (50:00:1f:e1:00:15:40:88) and Controller B, port 1 (50:00:1f:e1:00:15:40:8d). In the figure, the storage system host WWN is highlighted and the **Add Host >** button is used to place this storage system into the fabric.

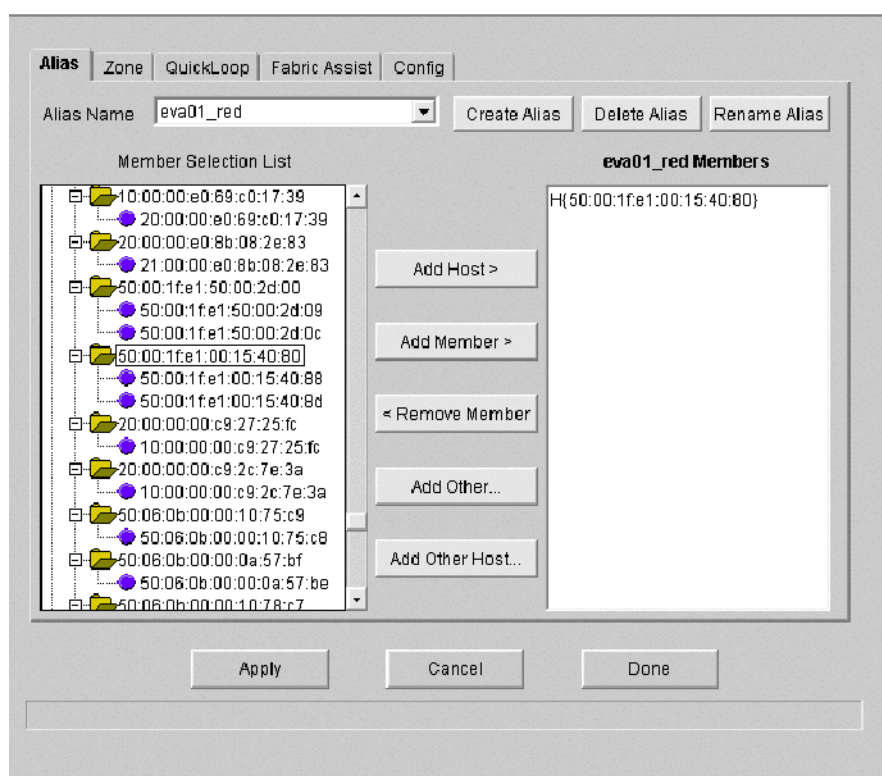


Figure 17: Example of host zoning with infrastructure switches

On the other fabric, the storage system WWN would display as 50:00:1f:e1:00:15:40:80, with Controller A, port 1 shown as 50:00:1f:e1:00:15:40:89 and Controller B, port 2 shown as 50:00:1f:e1:00:15:40:8c. The storage system WWN is highlighted, and the **Add Host** > button is used to zone by the host WWN.

Zoning with C- and M-series fabric switches

For best practice zoning recommendations with HP StorageWorks Edge and Director switches, and Cisco MDS switches, refer to the *HP StorageWorks SAN Design Reference Guide*.

HSG80 zoning recommendations

HSG80 controllers running Data Replication Manager (DRM) are not supported in a zone that shares an SMA with Continuous Access EVA. HSG80 controllers that are not running DRM are supported. Likewise, DRM and Continuous Access EVA can share intersite links.

With HSG80 controllers running in a Continuous Access EVA configuration, every SMA controls the HSG80 controller by issuing instructions through the command line interface (CLI). For Array Controller Software (ACS) versions earlier than V8.7, use switch zoning to include the HSG80 controllers with the active SMA. For ACS V8.7 or greater, HP recommends the use of Selective Management with the HSG Element Manager to allow CLI control only with the active SMA. Enabling Selective Management on the active SMA locks out CLI control on all other SMAs. The CLI commands `SET DISABLE_MANAGERS` and `SET ENABLE_MANAGERS` also provide the same functionality as Selective Management. See the *HP StorageWorks HSG80 Array Controller V8.7 Command Line Interface Reference Guide* for the use of the CLI commands.

To use Selective Management:

1. Log in to the HSG Element Manager on the active SMA.
2. Choose an HSG80 storage system and click **Controllers**. The Controller Properties page is displayed ([Figure 18](#)).
3. Click the **General** tab.
4. For Selective Management, choose **Exclusive** and click **Submit**.
5. Repeat steps 2 through 4 for all HSG80 storage systems.

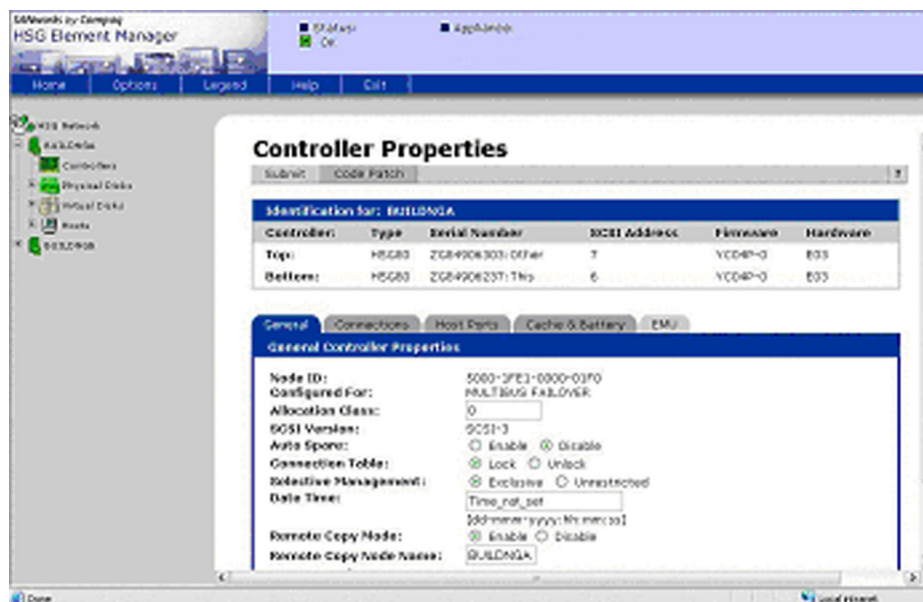


Figure 18: HSG Element Manager Controller Properties page

SMA-to-switch connections

The SMA comes equipped with two FCAs. Tag each end of your fiber optic cable to identify switch names, port numbers, appliance names, and so on. Connect the fiber optic cable so that each cable in the pair of FCAs goes to separate fabrics.

For procedures on synchronizing the time for multiple SMAs, see “Synchronizing time on the SMAs” on page 96.

Software configuration

This section describes various software configuration tasks that must be performed to begin running Continuous Access EVA. Procedures that are detailed in other documents or in online help are referenced.

Storage Management Appliance software setup

The SMA arrives with preinstalled software that allows you to add and run other applications. The two crucial applications to be added for Continuous Access EVA are HP StorageWorks Command View EVA and HP StorageWorks Continuous Access user interface.

Command View EVA

Command View EVA is a user interface that communicates with the HSV controllers to control and monitor the storage. Command View EVA maintains a database for each managed storage system and the database resides on that storage system. Configuration instructions for its use can be found in the *HP StorageWorks Command View EVA Getting Started Guide*.

HP recommends you use Command View EVA to:

- Initialize controllers
- Create and delete disk groups
- Create and delete hosts
- Create and delete host folders
- Create and delete Vdisks
- Create and delete Vdisk folders
- Create and delete snapshots
- Manage LUN presentations

Continuous Access user interface

The Continuous Access user interface is a Java applet created primarily for managing the various groupings of storage (copy sets, DR groups, and managed sets) in the Continuous Access EVA environment. For the application installation instructions, refer to the *HP StorageWorks Continuous Access EVA User Interface Installation Guide*. Detailed operational information is available with the Continuous Access user interface online help system.

Recommended uses for the Continuous Access user interface

HP recommends that you use the Continuous Access user interface to:

- Create and delete DR groups
- Change attributes for DR groups


- Add and remove members from DR groups
- Fail over, suspend, resume, and disable failsafe by DR group
- Create and delete managed sets
- Fail over, suspend, resume, and disable failsafe by managed set
- Fail over, suspend, resume, and disable by storage system
- Direct commands to the appropriate storage system regardless of selected storage system
- Actively monitor storage systems (copy, failsafe, logging, merging)
- Provide external notification via integration with the Notification Utility
- Create and restore configuration databases for managed sets and system folders.

Monitoring with the Continuous Access user interface

The Continuous Access user interface can be configured for event monitoring. An *event* is any reportable change in the storage system that is communicated by the HSV controllers. Examples of these events are:

- State changes in hardware
- State changes in a logical element, such as a Vdisk or DR group
- Procedure completions
- Environmental changes
- Operational changes

The procedure for setting up event monitoring is described in the *HP StorageWorks Continuous Access User Interface Installation Guide*. [Figure 19](#) shows the event notification flow. Events generated by the HSV controllers flow to the Command View EVA application on the SMA. Command View EVA sends the events as SNMP traps to the Continuous Access user interface and the Notification Utility. The Continuous Access user interface displays the events through an Internet browser. The Notification Utility can send information, warnings, errors, and failure notices as SNMP traps to other monitoring applications, or directly to e-mail addresses and pagers. See Appendix A for a listing of Continuous Access EVA events.



CXO8023B

Figure 19: Event notification flow

Configuring hosts

Be sure that you are using a supported Continuous Access EVA operating system. Supported operating systems are listed under “[Host operating systems](#)” on page 24, and supported versions of these operating systems are listed in the *HP StorageWorks Continuous Access EVA Design Reference Guide*.

Each operating system has a platform kit containing drivers and supporting documentation for use with the EVA. Also refer to any release notes for the platform-specific kit for any last-minute updates. The kit documentation will provide installation, upgrade, and testing information to ensure host connectivity to your fabrics.

Licensing

License keys must be obtained before your storage systems can be used for Continuous Access EVA or Business Copy EVA. To obtain the required keys, refer to the *HP StorageWorks Continuous Access EVA Getting Started Guide*. [Table 6](#) describes two types of license keys.

Table 6: License key types

Type	Description
Continuous Access EVA	Needed to unlock the Continuous Access features of the EVA. The license is based on replicated capacity, in Terabytes. One license per replicating array is required.
Business Copy EVA	Needed to unlock the snapshot and Snapclone features, if purchased. This license can be added any time before or after the system is initialized. The license is based on replicated capacity of the system connected to the controller, in Terabytes.

Initializing your storage systems

The storage system must be initialized before it can be used. This process binds the controllers together as an operational pair and establishes preliminary data structures on the disk array.

Initialization is performed through the use of Command View EVA. This procedure is documented in the *HP StorageWorks Command View EVA Getting Started Guide*.

Creating disk groups

Before creating a disk group, ensure that you have reviewed the planning considerations detailed in Chapter 3. A disk group is a set of physical disks from which storage pools are created. When your storage system is initialized, one default disk group is created. For the highest performance and availability, a disk group should only contain one disk drive model. If you are performing bidirectional replication, the disk groups should be symmetric with respect to the capacity on both storage systems.

Disk group creation is performed through the use of Command View EVA. The procedure for adding more disk groups is documented in the HP StorageWorks Command View EVA Online Help.

Creating host folders

An initial folder structure for a storage system is provided when you install Command View EVA. You can create additional subfolders to organize various components and to include hosts. For example, a folder can be used to store all hosts for a specific operating system. To organize a number of hosts into a category you wish to define, click the top-level host folder in Command View EVA, and then add a subfolder.

Creating hosts

The host connects to a fabric through an FCA and accesses storage through the controllers. Hosts contain a pair of FCA ports to connect with each fabric. Before a host can access any storage, it must be known to the storage system.

Hosts are created with Command View EVA while online and logged in to a fabric, so a path is defined between the Fibre Channel adapter and the storage system. In addition to assigning a host name, you must give the system an IP address for the host and the WWNs of the FCAs. Refer to the HP StorageWorks Command View EVA Online Help for information on creating hosts, modifying host properties, working with ports, and deleting hosts.

Note: Command View EVA allows a host definition to contain any number of ports. Continuous Access EVA requires that a host definition be limited to two ports—one on each fabric. This is due to a restriction that members of a DR group be preferred to the same controller and presented to the same host FCA pair.

Creating Vdisk folders

An initial folder structure for a storage system is provided when you install Command View EVA. You can create additional subfolders to organize various components to include Vdisks. To organize a number of Vdisks into a category you wish to define, click the top-level host folder in Command View EVA, and then add a subfolder.

Creating Vdisks

The Vdisk is a simulated disk drive created from a pool of disk group storage. You can assign a combination of characteristics to a Vdisk, such as a name, redundancy level, size, and other performance characteristics. Refer to the HP

StorageWorks Command View EVA Online Help for information on Vdisk creation using Command View EVA. The creation process may take several minutes or considerably longer, depending on the size of your Vdisk.

You are given the opportunity to select a preferred path during the creation of a Vdisk. This means that host I/O to a Vdisk will go to the controller you designate as preferred, as long as the paths to that controller are available. There are five possible preferred path settings:

- None
- Path A—Failover only
- Path B—Failover only
- Path A—Failover/failback
- Path B—Failover/failback

The default setting during Vdisk creation is None. If you are running an operating system that requires Secure Path software (HP-UX, IBM AIX, Linux, Novell NetWare, Sun Solaris, or Microsoft Windows) and want to designate a preferred path, use the “failover only” options. The failover/failback options are not supported with Secure Path. The two failover-only options allow the host to control when a Vdisk moves to a preferred path. For example, if Path A is preferred, and that path becomes unavailable, Path B is used. The host will then control the movement back to Path A when it becomes available later.

If you are running an operating system that does not require Secure Path software (OpenVMS or Tru64 UNIX), all the preferred path settings are supported. If failover and failback are desired, the “failover/failback” option should be used. This option allows the controller to manage the path movement of I/O to the Vdisk.

All operating systems require the entry of a logical unit number. You are given the opportunity to change the OS unit ID during the creation of a Vdisk. For OpenVMS and Tru64 UNIX, enter a unique OS unit ID for each Vdisk. The OS unit ID can be changed by clicking the Vdisk active member and selecting the Presentation tab.

Caution: All Vdisk members in a DR group must be preferred to the same controller and host FCA port pair because they must fail over together.

Installing Secure Path

Continuous Access EVA requires a multipath environment. Each host must have multipathing software installed, with the exception of HP OpenVMS and HP Tru64 UNIX (multipathing is built into these operating systems).

For the other operating systems requiring multipath capability, supported versions of HP StorageWorks Secure Path are required. Install Secure Path using the documentation that is provided with your platform-specific kit, in addition to any applicable release notes. Some operating systems may require special configurations when running Secure Path. For example, IBM AIX, and Linux with LifeKeeper Clusters, require a dual-path configuration that uses only one controller port on each controller of the pair.

Presenting Vdisks to hosts

Assigning a virtual disk to a host results in a host presentation. You may present a Vdisk to a host during the Vdisk creation process, or wait until a later time. However, the Vdisk must be presented to a host in order to use it in a DR group. The presentation is done with Command View EVA and is documented in the HP StorageWorks Command View EVA Online Help.

When Vdisks become members of a DR group, as described later in this chapter, all must have the same preferred path and host presentation.

Accessing the Continuous Access user interface

After you create hosts, Vdisks, and their presentations, and your hosts can access your Vdisks through Secure Path or multipathing, you are ready to create the Continuous Access EVA data structures (copy sets, DR groups, and so on). HP recommends that the Continuous Access user interface be used for these tasks. The procedures are described below and are also detailed in the Continuous Access user interface online help.

To access the Continuous Access user interface:

1. From a client computer, start your Internet browser and navigate to the SMA URL. The Device Home Page appears ([Figure 20](#)) and indicates successful connection to the SAN.

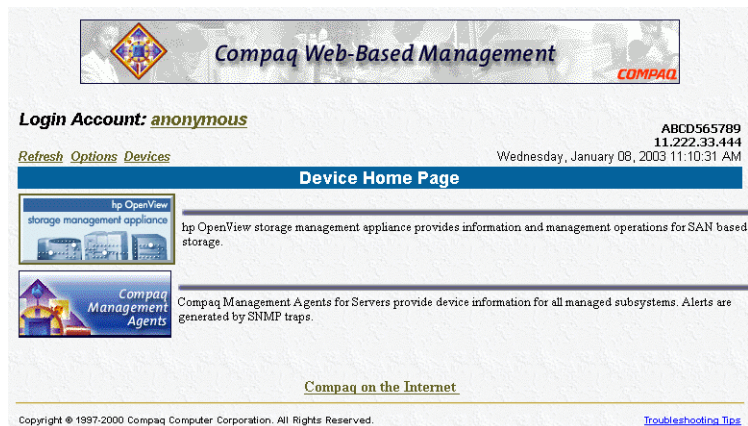


Figure 20: SMA Device Home page

2. Click the **hp OpenView storage management appliance** button. The SMA login window opens (Figure 21).

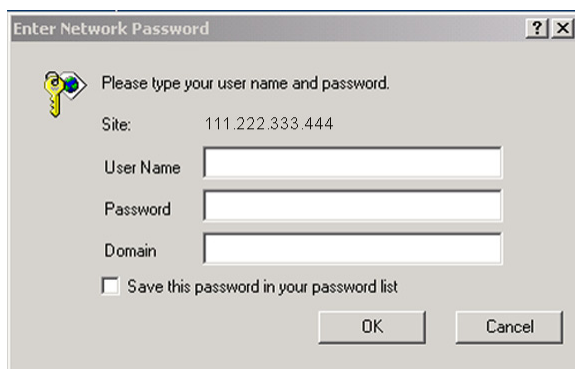


Figure 21: SMA login window

3. Log in to the SMA by entering a user name and password, and then click **OK**. The SMA software Home page is displayed. (Figure 22).

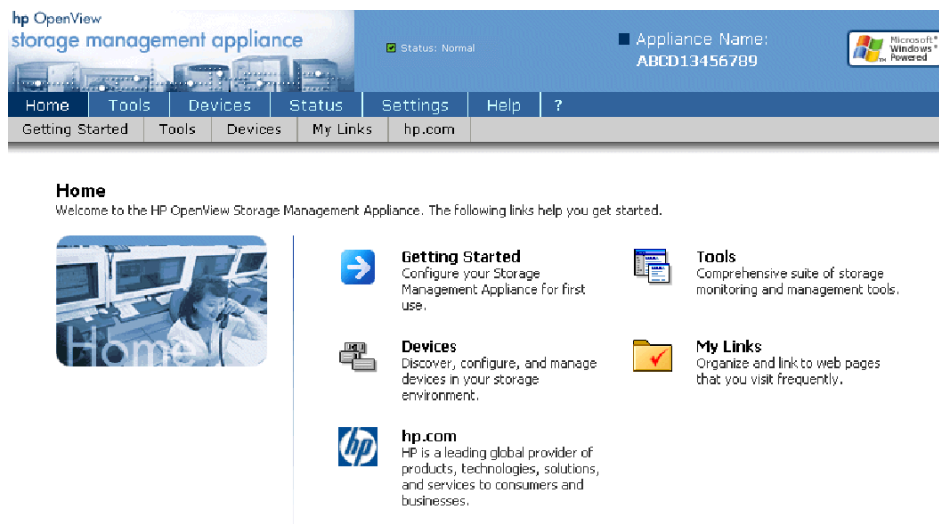


Figure 22: SMA software Home page

- Click the **Tools** icon. The Tools page is displayed (Figure 23).

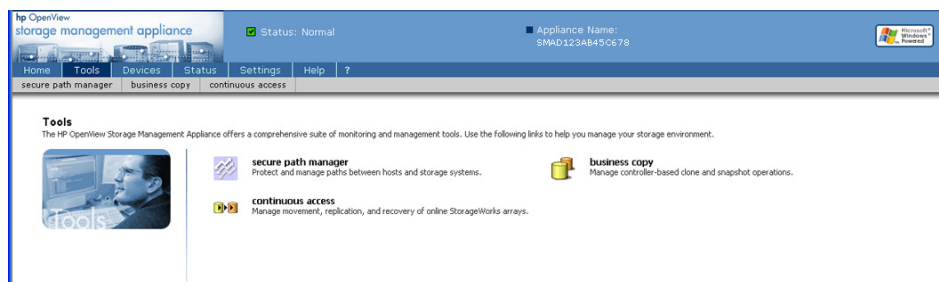


Figure 23: SMA Tools page

- Click the **continuous access** icon. The Continuous Access user interface main window opens (Figure 24).

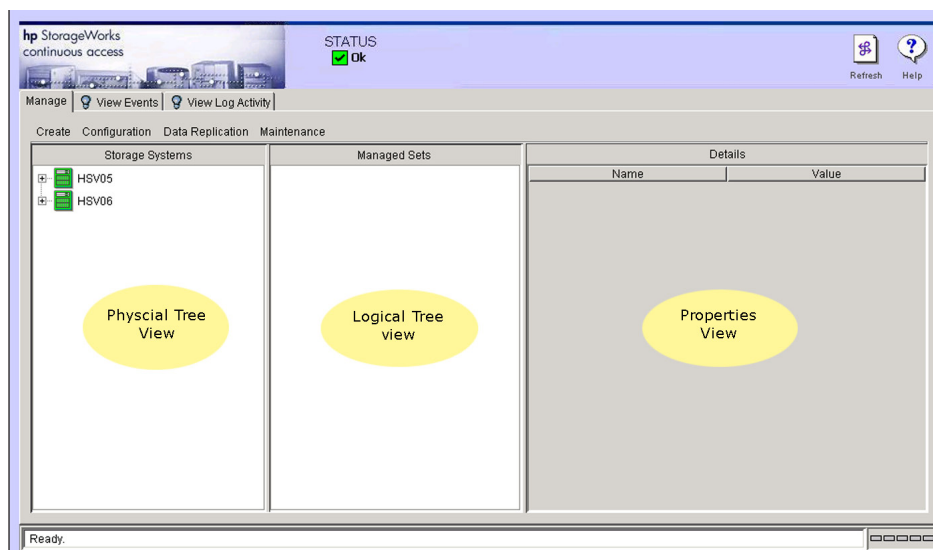


Figure 24: Continuous Access user interface main window

6. Click the **Refresh** icon and then choose **Rescan SAN**. The scanning process forces the Continuous Access user interface to identify all storage.

Creating copy sets and DR groups

The creation of a copy set results in the creation of a DR group if none exists. Afterward, copy sets may be created without the creation of more DR groups. Adding more Vdisks to an existing DR group will create a copy set based on that Vdisk. Placing copy sets into DR groups allows you to manage the replication behavior of them all as a group instead of individually.

To create a copy set:

1. Launch the Continuous Access user interface.
2. In the Storage System pane, choose the storage system folder that contains the Vdisks from which you want to create copy sets. See [Figure 25](#).

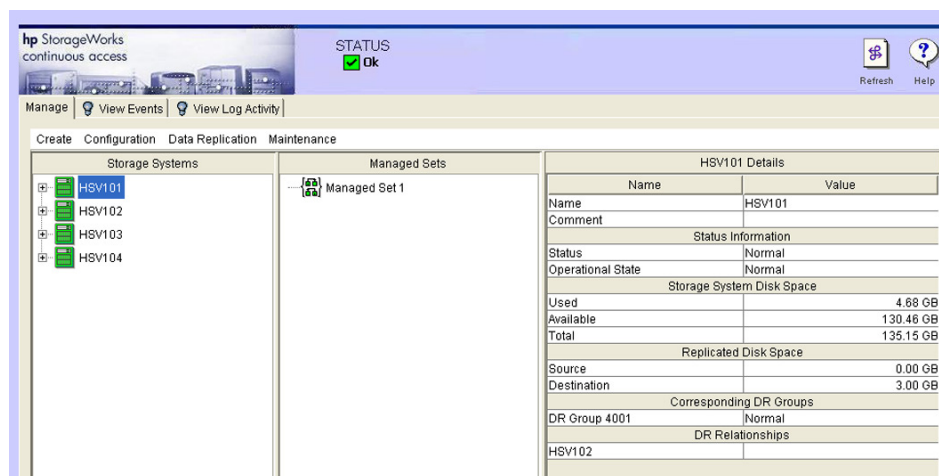


Figure 25: Storage selection on Continuous Access user interface main window

3. Click **Create > Copy Set**.

The Continuous Access user interface displays a window that asks if you want to put the new copy set into a new DR group (see [Figure 26](#)).




Figure 26: Add the new copy set window

4. You have three options. This step assumes you want to add the copy set to a new DR group.
 - If you do not want to create a DR group, click **Cancel** and proceed to step 7.
 - To add a copy set to an existing DR group, select **Existing**.
 - To add the copy set to a new DR group, click **New**. The Create a new DR Group window opens, as shown in [Figure 27](#).

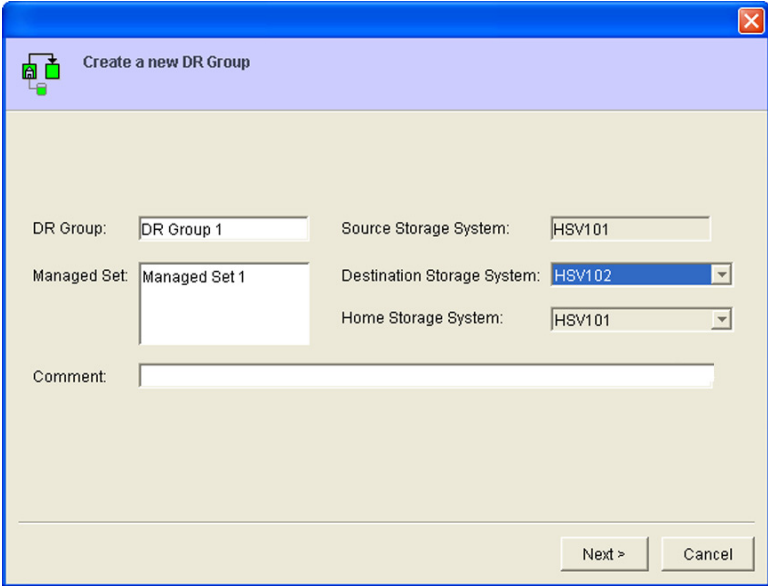


Figure 27 shows the 'Create a new DR Group' window. The window has a blue title bar with a close button. Below the title bar is a light blue header area with a small icon and the text 'Create a new DR Group'. The main area is light beige and contains several input fields: 'DR Group:' with a text box containing 'DR Group 1'; 'Source Storage System:' with a text box containing 'HSV101'; 'Managed Set:' with a dropdown menu showing 'Managed Set 1'; 'Destination Storage System:' with a dropdown menu showing 'HSV102'; 'Home Storage System:' with a dropdown menu showing 'HSV101'; and 'Comment:' with a text box. At the bottom right are 'Next >' and 'Cancel' buttons.

Figure 27: Create a new DR Group window

5. Complete the fields in the window and click **Next**.

Note: If the Vdisk being used for a new copy set does not appear in the Storage System pane, you must exit the Continuous Access user interface and present the Vdisk to a host by using Command View EVA. After returning to the Continuous Access user interface, perform a discovery.

6. Click **Next**, and the Create a new Copy Set window opens (see [Figure 28](#)).
7. Complete the fields in the window, and click **Finish**.

The new copy set is displayed in the Storage Systems pane. It is subordinate to the controller you selected and under the DR group you created or selected.

Figure 28: Create a new Copy Set window

When a copy set is created, storage is allocated from the selected disk group on the destination storage system, for the destination Vdisk. This creation process can take some time, being proportional to the capacity of the Vdisks. When the initial normalization completes, the copy set state is reported as normal because host I/O replication can begin, even though the storage allocation may still be in progress. Command View EVA and the Continuous Access user interface do not display that allocation is in progress at a copy set level, but the information can be retrieved at the Vdisk level, as described in the next step.

8. Verify that copy set storage allocation is complete by going to the Command View EVA Vdisk Active Member Properties page (Figure 29) for the destination copy of the Vdisk. Allocation is in progress when the **Capacity Used** value is less than the **Capacity Req** total. Allocation is complete when both values are the same. A failover of the DR group cannot be performed and disaster tolerance is not achieved until both the allocation and full copy data copy are complete.



Caution: After a copy set is created, the destination Vdisk must be presented to a host by using Command View EVA. If a failover occurs, the host will not be able to access the destination Vdisk until this step is performed.




Figure 29: Storage allocation check on Vdisk Active Member Properties page

Presenting a copy set to a destination host

A copy set must be presented to a host at the destination storage system so the host can access it upon failover. This procedure is performed with Command View EVA and is described in the HP StorageWorks Command View EVA Online Help.

Specifying disk group membership for a log

During the creation of a DR group, a 139-MB log is automatically created for the source and destination storage systems. Placement of the log into a disk group on each system is based on the amount of free space in the disk group, the number of other logs in each disk group, the potential size of each existing log, and the potential size of the new log. It is possible that the disk group automatically selected for the log may not be the same disk group where the DR group resides. For example, disk groups containing near-online FATA drives are automatically selected for the location of the log. If near-online drives are not present, disk groups with the largest amount of average free log space are chosen. The location of the log is rarely an issue, but some users may want to control the placement of a log into a specific disk group to separate the log from the data.

By default, a log is created in a near-online disk group. You cannot force a log into an online disk group if a near-online disk group exists. This restriction applies to both the source and destination storage systems.

If possible, creating DR groups and specifying log disk location should be done before other non-DR Vdisks are created. Doing this ensures sufficient available space in the disk groups for a log.

Provided there is enough available space in the destination storage system to create the destination Vdisks and log disk, use the following procedure to manually specify the disk group where a log will be created:

1. Create Vraid0 Vdisks on the destination storage system in all disk groups except the one that will be used for the destination Vdisk and log, so that all available space is filled. These Vdisks are temporary and will not be used for data storage.
2. Click each of the disk groups in turn to see the available space on the destination storage system. Each disk group should report zero space available except the disk group where the destination Vdisk and log will be created.
3. Wait until the Vdisks have finished allocating space and are in a normal state.
4. Create Vraid0 Vdisks on the source storage system in all disk groups except the one that will be used for the log, so that all available space is filled. These Vdisks are temporary and will not be used for data storage.
5. Click each of the disk groups in turn to see the available space on the source storage system. Each disk group should report zero space available except the disk group where the log will be created.
6. Wait until the Vdisks have finished allocating space and are in a normal state.
7. Create the DR group. On the DR group creation page, use the Destination Storage System drop-down list to specify the disk group on the destination storage system where you want the destination Vdisk and log created. Do not allow the option to automatically select.
8. Delete the Vraid0 Vdisks that were created in step 1, and then delete those created in step 4.

The log disk is now located in the disk group you selected. The disk group membership of the log is not visible to Command View EVA or to the Continuous Access user interface.

Deleting or detaching copy sets

Copy sets can be either detached or permanently deleted from the destination storage system. The following procedure describes both options.

1. Choose the icon of the copy set to be deleted or detached.
2. Click **Configuration > Delete**.

A dialog box displays with three button options: **Delete**, **Detach**, and **Cancel**. The delete option permanently deletes the source and destination Vdisks making up the copy set. The detach option deletes the copy set but leaves the destination Vdisk intact.

3. Click the appropriate option.

Deleting DR groups

The deletion of a DR group will delete all copy set members in the group. In addition, this operation automatically deletes the destination Vdisks if the source-destination connection is operational. If the source-destination connection is down, this operation deletes the DR group, but destination Vdisks remain. When the connection is re-established, the destination Vdisks have to be manually removed. To delete a DR group:

1. Choose the icon of the DR group to be deleted.
2. Click **Configure > Delete**.
3. Select **Yes**.

Creating managed sets

A managed set is a folder created to hold DR groups. One or more DR groups can be combined to create a managed set, but there are no DR groups in the managed set when it is created. Perform the following steps to create a managed set:

1. Choose **Create > Managed Sets**. The Edit or create a Managed Set window opens ([Figure 30](#)).
2. Enter the Managed Set Name, and specify any optional comments. Managed set names are limited to 64 characters and are case-sensitive.
3. Click **Finish**.

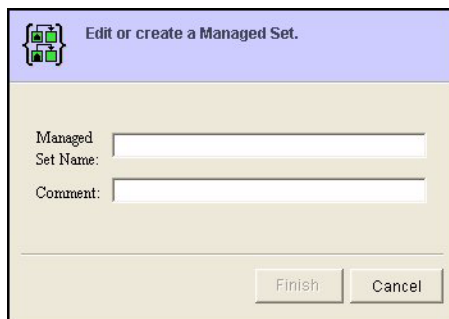


Figure 30: Edit or Create a Managed Set window

Editing a managed set

Only the name or comments of a managed set can be edited. To edit these:

1. Choose the icon of the managed set.
2. Choose **Configuration > Edit** on the **Manage** tab. The Edit or create a Managed Set window opens.
3. Make the appropriate changes.
4. Click **Finish**.

Adding a DR group to a managed set

The following steps describe how to add DR groups to a managed set:

1. Choose a DR group icon in the Storage Systems pane.
2. Choose **Configuration > Edit** on the **Manage** tab.
3. In the **Managed Set** box within the Create a new DR Group window, choose a managed set. If you want to include the DR group in more than one managed set, then **Ctrl-click** the managed set.

Note: To cancel your choice, press **Ctrl** while you click the appropriate managed set on the list.

4. Click **Next**, and then click **Finish**.

Removing a DR group from a managed set

To remove a DR group from a managed set:

1. Choose the icon of the DR group you want to remove in the Storage Systems pane.
2. Click **Configuration > Edit** on the **Manage** tab. The Edit an Existing DR Group window opens.
3. In the Managed Set field, click on the managed set from which you want the DR group removed.

If you want to remove the DR group from more than one managed set, **Ctrl-click** the managed set.

4. Click **Finish**.

Deleting a managed set

Perform the following steps to delete a managed set:

1. Click the appropriate managed set in the Continuous Access user interface main window.
2. Click **Configuration > Delete** on the **Manage** tab. A dialog box displays **Yes** and **No** options.
3. Click **Yes**.

Backing up configuration information

There are two procedures you should perform to back up specific information to help re-create a Continuous Access EVA configuration. Each procedure uses a different technique and captures data sets used to help recover a saved configuration.

The first procedure backs up the information unique to the Continuous Access user interface environment and is done through the use of the Continuous Access user interface. This configuration data is important when moving management control to another SMA, so it is discussed in reference to that use in the section “Saving your Continuous Access EVA storage configuration” on page 89.

The second procedure backs up the configuration data about an array and is done using the Storage System Scripting Utility (SSSU) to survey specified storage systems and create configuration recovery files that can be saved to a directory on a platform server.

Note: The SSSU is not available on Novell NetWare hosts or on the SMA.

Using the SSSU `Capture Configuration` command, five separate scripts are run that will append a user-defined configuration name to the file with a name in the form *UserName_StepX.txt*, where StepX is one of these five configuration text files:

Step1A—This step captures the data needed to re-create the storage system itself, disk groups, hosts, Vdisks that are not used for data replication (either source or destination), and LUNS for the disks created.

Step1B—This step captures the data needed to re-create all source Vdisks used in DR groups on this storage system. However, the data captured by this step is not currently used in any recovery procedures.

Step1C—This step captures the data needed to present all source Vdisks (creates LUNs) used for DR groups to their hosts. However, the data captured by this step is not currently used in any recovery procedures.

Step2—This step will capture the data needed to re-create all data replication specific configuration information only, and only that DR-specific information for which the storage system is the source. This consists of source DR groups and their copy set members only.

Step3—This step will capture the data needed to create an SSSU script that will again present all remote Vdisks to their hosts.

Example:

You have an SMA with an IP address of 111.222.333.444. You wish to back up your configuration for a storage system named HSV01 with the SSSU. Follow these steps:

1. Run the SSSU executable (*SSSU.exe*) to get a command prompt.
2. At the command prompt, log into the SMA using your user name and password. For example:

```
select manager 111.222.333.444 username=user1 password=admin
```

3. Select the storage system whose configuration you want to save. Use the following command with your storage system name (HSV01 is the storage system name used in this example):

```
select system HSV01
```

The command line prompt changes to reflect your storage system is selected.

4. At the SSSU command prompt, enter the `Capture Configuration` command along with a path and file name where you want the configuration text files to reside. For example, to copy these files to a folder called `storage_systems\hsv01`, use the command:

```
capture configuration c:\storage_systems\hsv01.txt
```

You will see confirmation messages on the screen that each step was successfully saved.

A current copy of the data that defines the storage system configuration has been saved. If configuration changes are made later to your storage system, then you will have to rerun this procedure. Procedures that describe how to recover a configuration using the SSSU are part of the failover procedures titled “Return operations to replaced new storage hardware” on page 117.

Storage Management Appliance Procedures

5

This chapter describes how to use one or more Storage Management Appliances (SMAs) with Continuous Access EVA. The following topics are discussed:

- [Considerations for managing storage with multiple SMAs](#), page 88
- [Saving your Continuous Access EVA storage configuration](#), page 89
- [Stopping SMA applications](#), page 90
- [Restarting SMA applications](#), page 92
- [Moving storage management to another SMA](#), page 93
- [Synchronizing time on the SMAs](#), page 96
 - [Setting a storage system time to an SMA](#), page 96
 - [Synchronizing SMAs to an NTP server](#), page 98
- [Enabling management on an SMA when HSG80 controllers are present](#), page 99
 - [The managing SMA is still accessible](#), page 99
 - [The managing SMA is disabled](#), page 100

For a general discussion describing the use of multiple SMAs in a SAN, refer to *HP OpenView Storage Management Appliance Software Using Multiple Storage Management Appliances in a SAN Application Notes*.

Considerations for managing storage with multiple SMAs

A storage system can only have one SMA managing it. Multiple SMAs can be used to divide up management responsibilities so that each SMA controls different storage systems. Any SMA residing in the management zone with your storage systems has the ability to acquire the control or management of any storage system from another SMA.

Here are some points to consider when using multiple SMAs:

- Synchronize the clocks of the SMAs, taking time zones into account. This is important for event and log time stamps. (See “Synchronizing time on the SMAs” on page 96.)
- Decide which SMA will manage each storage system and which SMA can be used to take control of that storage system, if needed. (See “Moving storage management to another SMA” on page 93.) A single SMA must manage both storage systems when creating DR groups. Afterwards, separate SMAs may be used to monitor the state of the storage systems.
- Realize what applications are installed and running on each SMA. Should another SMA be needed to control a storage system, the SMA will need the same applications running to allow the same functionality. Applications can also be stopped and started on an SMA. (See “Stopping SMA applications” on page 90 and “Restarting SMA applications” on page 92.)
- Document and save your Continuous Access EVA storage configurations for all SMAs, especially relating to DR groups, system folders, and managed sets. Configurations can be saved with the Continuous Access user interface backup utility as described in “Saving your Continuous Access EVA storage configuration” on page 89.

The storage configuration on an SMA may change after acquiring a new management role. You can pre-configure the SMA to prepare for that role, then back up the configuration to be used when needed. The SMA can perform one management role until a backed up configuration is restored to allow it to assume another role.

- Practice SMA management role changes regularly. This allows you to verify that clock times, installed applications, and storage configurations are correct.

Saving your Continuous Access EVA storage configuration

The following procedure describes how to save your Continuous Access EVA storage configuration using the Continuous Access user interface.

1. Create a folder on a host (preferably not the SMA) where you want to save your Continuous Access EVA configuration files.
2. Log in to the Continuous Access user interface on the SMA. (See “Accessing the Continuous Access user interface” on page 72.)
3. On the Continuous Access user interface main window, click **Maintenance** and then click **Backup Database**. A Select Location window opens (Figure 31) from the host where you are browsing, allowing you to select a location for the backup information.

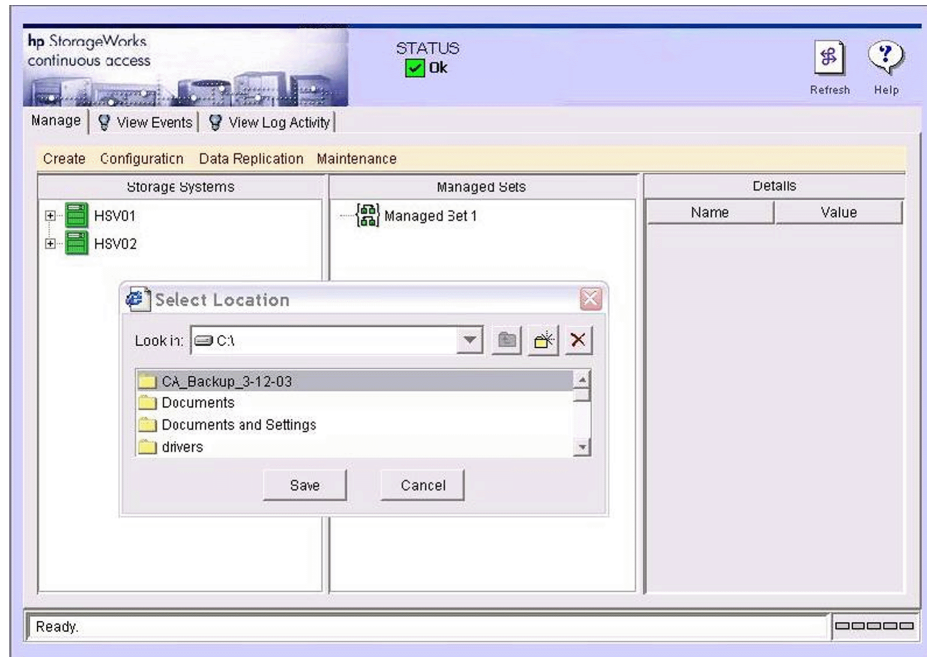


Figure 31: Continuous Access user interface Select Location window

4. Select the folder you created for the Continuous Access EVA configuration backup file and click **Save**. A Download window opens ([Figure 32](#)).




Figure 32: Continuous Access user interface Download window

Procedures for restoring from the backup database are described in “Moving storage management to another SMA” on page 93.

Note: If you elected to save your backup configuration to a folder that already contains configuration files, you will see a Destroy Files message window. If you want to keep the old configuration files in that folder, click **No** and create a new folder for the current backup files. Otherwise, click **Yes** to overwrite the old backup configuration files.

Stopping SMA applications

Use the following procedure to stop applications running on an SMA.

1. Log in to the SMA. The SMA software Home page is displayed ([Figure 22](#)).
2. Click **Settings** on the primary navigation bar. The SMA Settings page is displayed ([Figure 33](#)).
3. Click **Manage tools**. The Manage Tools page is displayed ([Figure 34](#)) and lists applications installed on the SMA.

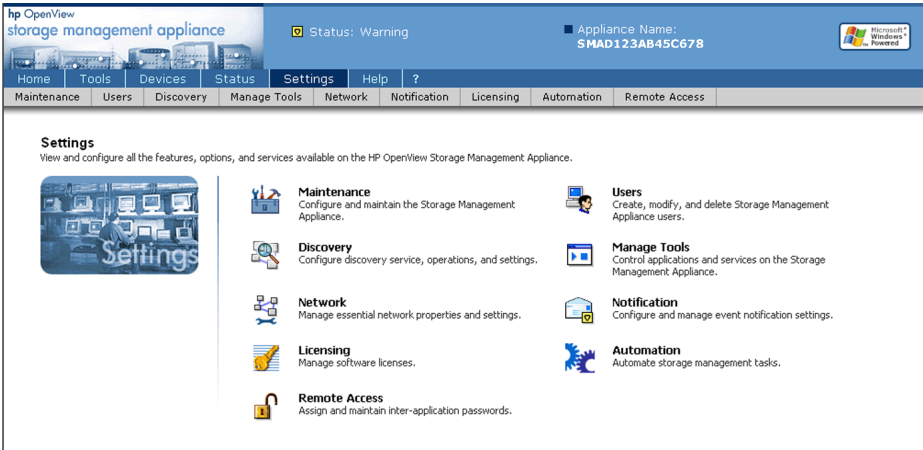


Figure 33: SMA Settings page

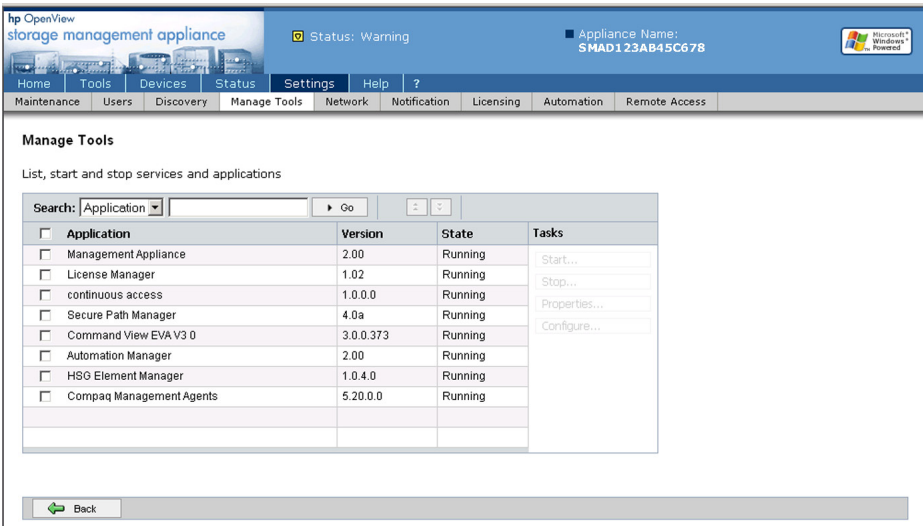


Figure 34: SMA Manage Tools page

4. Choose the check boxes of the applications to be stopped. The **Stop** and **Properties** tabs become visible under the **Tasks** section (Figure 35).
5. Click **Stop**. The state of the applications changes to show they were stopped.




Figure 35: Selected applications to stop on the Manage Tools page

Restarting SMA applications

Use the following procedure to restart applications on the SMA.

1. Log in to the SMA and navigate to the Manage Tools page as described in “Stopping SMA applications” on page 90.
2. Choose the check boxes of the applications to be restarted. The **Start** and **Properties** tabs become visible under the **Tasks** sections (Figure 36).




Figure 36: Selected applications to restart on the Manage Tools page

3. Click **Start**. The state of the applications changes to show they are running.

Moving storage management to another SMA

The following procedure describes how to move management of EVA storage systems from one SMA to another SMA. If your configuration contains HSG80 controllers, see “Enabling management on an SMA when HSG80 controllers are present” on page 99 for additional procedures that must be performed.

1. Log in to Command View EVA on the SMA that you want to use to control your storage system.
2. Choose **Discover**, then **OK**. The storage icons displayed in the navigation pane are gray to indicate that another SMA is managing the storage.
3. Click a storage icon in the navigation pane. A page similar to [Figure 37](#) is displayed to inform you that another instance of Command View EVA is managing the storage systems.

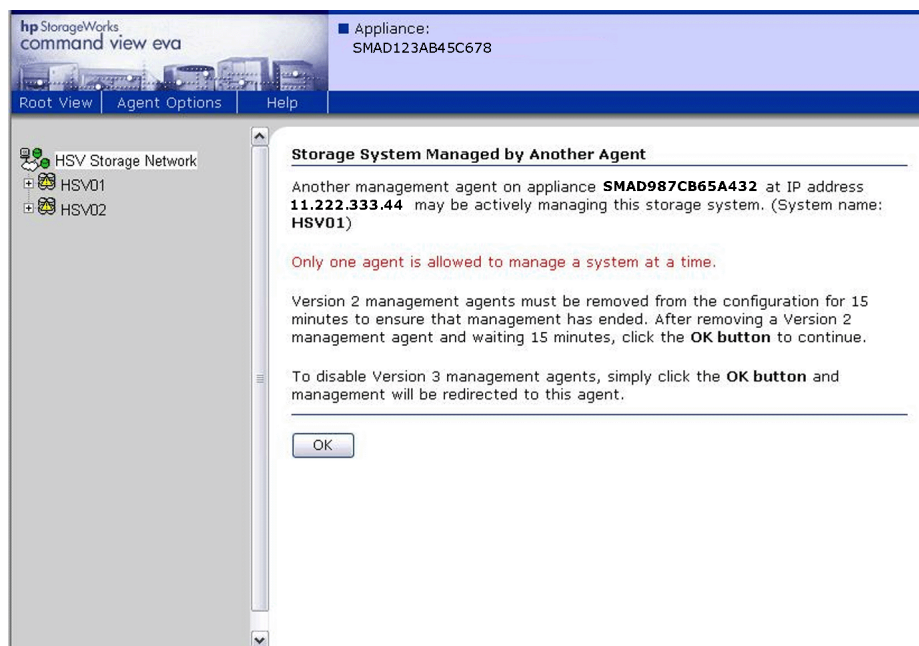


Figure 37: Storage System Managed by Another Agent page

4. Click **OK**. A message similar to Figure 38 indicates that you are about to assume control of the storage with another SMA.




Figure 38: Assuming management of your storage message

5. Click **OK**. The storage icon in the navigation pane becomes green to indicate that the SMA you are logged in to has control of this storage system. The color change to green may take a little while to occur.

6. Repeat steps 3 through 5 for the remaining gray storage icons for which you want to take control.
7. Launch the Continuous Access user interface. See “Accessing the Continuous Access User Interface” on page 71.
8. Choose one of the following:
 - a. If you have a backup of your configuration that you wish to restore, click **Maintenance > Restore Database**. A Select Location window opens ([Figure 31](#)). Choose the folder containing your Continuous Access EVA backup files, select the desired backup file, click **Restore**, click **Yes** in the Restore window, and then click **OK**.
 - b. If you do not have a configuration backup to restore, create your managed sets and system folders using configuration information you have documented.
9. Display your restored configuration.
 - a. If you did not restore configuration from a backup, create your managed sets.
 - b. If you restored a configuration backup, or if the storage configuration is not seen or correctly displayed, click the **Refresh** icon. A Discover window opens. Select **Rescan SAN**. After a rescan is performed, the storage icons turn green to indicate a normal condition.
 - c. If your storage configuration is correctly displayed but the states are incorrect (as indicated by color), click the **Refresh** icon. A Discover window opens. Select **Refresh Display**. The status of the storage systems is verified and their icons are green to indicate a normal condition.

Synchronizing time on the SMAs

It is important that the time clocks of your SMAs be synchronized as close as possible, keeping time zone differences in mind. This allows you to view event times in log files from multiple SMAs with a consistent time reference. This is especially important when comparing event logs of both storage systems in a replicating relationship. There are two layers to the synchronization process:

- Setting a storage system to the time of a particular SMA
- Synchronizing SMAs to a Network Time Protocol (NTP) server

Setting a storage system time to an SMA

The following procedure describes how to set the time of a storage system to an individual SMA:

1. Log in to the SMA managing the desired storage system.
2. Log in to Command View EVA.
3. Choose the desired storage system. The Initialized Storage System Properties page is displayed ([Figure 39](#)).

The screenshot shows the HP StorageWorks Command View EVA web interface. The top navigation bar includes 'Root View', 'Agent Options', and 'Help'. The left sidebar shows a tree view of the 'HSV Storage Network' with nodes 'HSV07' and 'HSV08'. The main content area displays the 'Initialized Storage System Properties' page for 'HSV07'. The page has a title bar with 'Appliance: SMAD 111.222.333.444'. Below the title bar are tabs: 'Save changes', 'Set options', 'View events', 'Uninitialize', and a help icon. Below these are buttons: 'Code load' and 'Shut down'. The main content is divided into two columns: 'Identification' and 'Condition/State'. The 'Identification' column contains fields for 'Name' (HSV07), 'Node World Wide Name' (5000-1FE1-0015-31E0), and 'UUID' (6005-08b4-0001-446f-0003-9000-0382-0000). The 'Condition/State' column contains 'Operational state' (Good (Initialized)) and 'System' (Type: HSV110, Version: 3010, Console LUN ID: 0).

Identification		Condition/State	
Name:	HSV07	Operational state:	<input checked="" type="checkbox"/> Good (Initialized)
Node World Wide Name:	5000-1FE1-0015-31E0		
UUID:	6005-08b4-0001-446f-0003-9000-0382-0000		
		System	
		Type:	HSV110
		Version:	3010
		Console LUN ID:	0

Figure 39: Initialized Storage System Properties page

4. Click the **Set Options** tab. The System Options page is displayed ([Figure 40](#)).

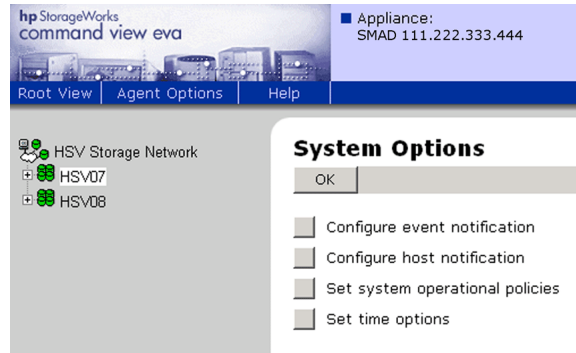


Figure 40: System Options page

5. Choose the **Set time options** box. Various options for setting time are displayed on the Set System Time page (Figure 41).




Figure 41: Set System Time page

6. Click **Use management appliance date/time**.

7. Click **Re-sync controller time with the SAN management time**.
8. Click the **Save changes** tab.

Synchronizing SMAs to an NTP server

To synchronize the time of all your SMAs to a single NTP server:

1. Log in to the SMA whose time is to be set.
2. On the SMA Home page, click **Settings**.
3. Click **Maintenance**.
4. Click **Date/Time**.
5. On the Date and Time Settings page, select the **Synchronize the Appliance Date and Time with this SNTP Server** box (see [Figure 42](#)).

hp OpenView
storage management appliance

Status: Informational Appliance Name:

Home Tools Devices Status Settings Help ?

Maintenance Users Discovery Groups Manage Tools Network Licensing Notification Automation

Date and Time Settings

Date:

Time:

Time zone:

☒ Automatically adjust clock for daylight saving changes

Changing the Storage Management Appliance date and time does not affect the date and time on your computer.

☒ Synchronize the Appliance Date and Time settings with this SNTP Time Server

SNTP Server (local or web):

Synchronization Period:

Figure 42: Date and Time Setting page

6. Enter an IP address for your SNTP server in the **SNTP Server (local or web)** field.
7. In the **Synchronization Period** drop-down box, select how often you want the time resynchronized.

Enabling management on an SMA when HSG80 controllers are present

If your Continuous Access EVA configuration contains HSG80 controllers, additional steps are necessary to move storage management to another SMA. These steps, performed in addition to those involving the movement of storage management to another SMA with Command View EVA and the Continuous Access user interface, allow the SMA to manage storage presented by HSG80 controllers. Separate procedures are necessary, depending on the condition of an SMA before moving storage management to another SMA. These conditions are the following:

- The managing SMA is still accessible.
- The managing SMA has become disabled, as would happen if the intersite links are down or the site is down where the managing SMA resides.

Note: HSG80 controllers running Data Replication Manager (DRM) cannot be in the same zone with any SMA. Refer to “HSG80 Zoning Recommendations” on page 62.

The managing SMA is still accessible

Perform the following steps for HSG80 controlled storage:

1. Log in to the HSG Element Manager on the managing SMA.
2. Select an HSG80 storage system and click **Controllers**. The Controller Properties page is displayed ([Figure 18](#)).
3. Choose the **General** tab.
4. For Selective Management, choose the **Unrestricted** option.
5. Choose the **Submit** tab.
6. Repeat steps 2 through 5 for all HSG80 storage systems.
7. Log in to the HSG Element Manager on the other SMA.
8. Choose an HSG80 storage system and choose **Controllers**.

Note: An HSG80 storage system icon becomes blue to indicate a CLI lockout condition. If this occurs, select the HSG80 storage system icon and click the **Rescan** button. When the rescan finishes, exit out of the System Properties window and complete step 7.

9. Choose the **General** tab.
10. For Selective Management, select the **Exclusive** option.
11. Choose the **Submit** tab.
12. Repeat steps 8 through 11 for all HSG80 storage systems.

The managing SMA is disabled

If the managing SMA is disabled, CLI commands cannot be performed from the other SMA except through the use of a serial maintenance port on the HSG80 port or through the use of StorageWorks Command Console (SWCC).

Use the following procedure to enable the other SMA to manage HSG80 storage if the managing SMA in your Continuous Access EVA configuration is disabled:

1. From the HSG serial port or the StorageWorks Command Console (SWCC), type `SHOW CONNECTIONS` and record the connection name of the standby SMA. If both Fibre Channel adapters (FCAs) on the standby SMA are installed, then record both connection names.
2. Type `SET DISABLE_MANAGERS=ALL` to disable CLI commands on the failed active SMA.
3. Type `SET ENABLE_MANAGERS=alternateSMAconnectionname`. To enable both FCAs, type:

```
SET ENABLE_MANAGERS=alternateSMAname1,  
alternateSMA name2.
```

See the *HP StorageWorks HSG80 V8.7 Array Controller Command Line Interface Reference Guide* for more information about the CLI command.

4. Type `SHOW MANAGERS` to verify changes.
5. Repeat steps 1 through 4 on any other HSG80 storage systems.
6. On the newly active SMA, log in to the HSG Element Manager.
7. Select an HSG80 storage system and click **Controllers**.

Note: An HSG80 storage system icon becomes blue to indicate a CLI lockout condition. If this occurs, select the HSG80 storage system icon and click the **Rescan** button. When the rescan finishes, exit out of the System Properties window, and log back into the HSG Element Manager.

8. Verify that all have selective management on the Controller Properties page set to **exclusive**.

Recovery

6

This chapter provides recovery information for performing failovers, resumption of operation after encountering a failsafe-locked condition, and when encountering a disk group failure. Procedures use the HP StorageWorks Continuous Access user interface, when applicable. Seven scenarios are provided that cover the majority of situations likely to be encountered, along with detailed procedures for handling each scenario.

The following major topics are discussed:

- [Planning for a disaster](#), page 104
- [Failsafe and normal modes](#), page 105
- [Throttling of merge I/O after logging](#), page 106
- [Failover defined](#), page 106
- [The Continuous Access user interface](#), page 110
 - [Continuous Access user interface icons](#), page 110
 - [Data replication using the Continuous Access user interface](#), page 113
- [Possible event scenarios](#), page 115
- [Failover and recovery procedures](#), page 119
 - [Planned failover](#), page 119
 - [Unplanned failover](#), page 127
 - [Resumption of operations if unable to access destination while source in failsafe-locked state \(extended period of time\)](#), page 132
 - [Return operations to home storage system](#), page 134
 - [Return operations to replaced new storage hardware](#), page 134
 - [Recovering from a disk group hardware failure](#), page 142

Planning for a disaster

When a disaster occurs at one of your storage sites, your immediate priority becomes getting your data back online in the shortest amount of time. There are many actions that can be planned ahead of time to minimize the downtime brought on by a disaster:

- Operating with a supported disaster-tolerant Continuous Access EVA configuration is of primary importance. Make sure there are two fabrics with at least one intersite link per fabric.
- Make sure your controllers are cabled in the supported “cross-cabling” configuration to your fabrics (see [Figure 16](#)).
- Have at least one Storage Management Appliance (SMA) available at every site in case of an appliance or communication failure.
- Verify that each destination Vdisk within a DR group has been presented to a host. This allows the host access to the Vdisk immediately after a failover.
- Make sure local and remote hosts are installed with the latest patches, virus protection, EVA platform kits, and Secure Path versions for that operating system.
- Keep your configuration current and documented at all sites. Install the latest versions of VCS firmware, SMA software, Command View EVA, and the Continuous Access user interface.
- Keep a record of your Vdisks and DR groups. An easy way to do this is through the use of the `capture` command with the Storage System Scripting Utility (SSSU) after significant changes or at scheduled intervals (see “Backing up configuration information” on page 84).
- Keep the Continuous Access user interface on every SMA synchronized with any configuration changes. After changes are made using one SMA, log on to your secondary SMA, launch Command View EVA, and choose each storage system to “take control” with the secondary SMA. Then launch the Continuous Access user interface and select **Refresh > Discover** to have the changes recognized. Repeat this on every SMA. This synchronization of the Continuous Access user interface prepares the standby SMAs for activation in the event of a disaster.
- Back up the Continuous Access user interface database for any configuration that the SMA may use. These databases contain managed set and system folder information that can be quickly restored when an SMA changes its role.

For example, two SMAs may control different storage systems at one time, but each SMA can have a database created with all the storage systems under its control in case the other one becomes inoperative.

A well prepared recovery plan can become worthless without practice. Make sure everyone involved in your storage administration practices for disaster recovery. Try out different failure scenarios. Make decisions ahead of time as to what scenarios lead to failover. For example, if a controller fails, is it more important not to disrupt processing by doing a planned failover, or to not be at risk for a second controller failure that will result in an unplanned failover? In the case of multiple sites, which site has precedence for troubleshooting?

Scheduling periodic times to practice disaster recovery can be a good occasion to verify that your records are up-to-date and that all required patches are installed.

Failsafe and normal modes

Failsafe is a DR group mode that can be set to stop I/O to the DR group if any member of the DR group (source or destination) becomes inaccessible. Replication goes into an inoperative state called *failsafe-locked*. No logging occurs, because write I/O is not being accepted by the active storage system. The reason for choosing this error mode is to ensure that the data at your source storage system exactly matches the data at your destination storage system.

For example, a source DR group has been set to failsafe-enabled. If the destination becomes inaccessible, then the source DR group becomes failsafe-locked. To continue processing without a destination DR group, you would have to disable the failsafe error condition from the source DR group and operate in normal mode.

If you have a source DR group that is failsafe-enabled, and it becomes failsafe-locked, you can fail over to the destination storage system and regain access to your destination DR groups from the host at your alternate site. When your original source becomes accessible again, a merge or full copy will occur. If the new source has failsafe-enabled DR groups, they will continue to operate in that mode when the connection is restored.

Several of the procedures in this chapter direct you to transition between failsafe-locked mode and normal mode so that you can resume processing at the source storage system. Changing between failsafe mode and normal mode does not constitute an actual failover event. It is a resumption of operation.

Throttling of merge I/O after logging

When I/O has been halted, DR groups not in failsafe mode will automatically resume replication when links to the remote storage systems are restored. If there are dozens of DR groups with large logs, they will compete for bandwidth as they try to synchronize simultaneously.

By suspending the merging or copying of non-critical DR groups, the controllers will merge only the most critical data first, allowing this data to be synchronized and become accessible before the less important data. As the more important groups finish merging, you can resume the I/O of the groups that were suspended. This concentration or channeling of I/O to specific groups by the use of suspend and resume commands is called *throttling I/O*.

Failover defined

Generically speaking, *failover* is the substitution of a functionally equivalent device or component for a failed one. Failover can take several forms with Continuous Access EVA:

- *Controller failover* is the process that takes place when one controller in a pair assumes the workload of a failed or redirected companion controller in the same cabinet.
- *DR group, managed set, or storage system failover* is an operation to reverse the replicating direction of the DR group, managed set, or storage system to its partner. This reversal is possible because all data generated at a source storage system has been replicated to a destination storage system, in readiness for such a situation.
- *Fabric or path failover* is the act of transferring I/O operations from one fabric or path to another.

This chapter discusses the failover of DR groups, managed sets, and storage systems. It does not discuss controller failover within a cabinet, or path or fabric failover, because redundancy is assumed.

The failover method used with DR groups, managed sets, or storage systems is determined by the severity of the failure or the reason for the failover. A *planned failover* can be used for situations such as an anticipated power disruption, scheduled equipment maintenance at the source storage system, or the need to transfer operations to another storage system. An *unplanned failover* is used for events such as multiple controller failures, multiple host failures, or an unplanned power outage at the source storage system.

If the source storage system fails, or if you are planning downtime with the source storage system, you must decide whether to perform a failover to the destination storage system. Always verify that all components at the destination storage system are operational before you begin a failover.

When you perform a failover, the destination storage system assumes the role of the source and becomes the active storage system. It remains the source storage system until you fail over to another system. By transferring control of system operation to the destination storage system, you can ensure minimal interruption of data access after a failure.

Note: When you perform a failover for a DR group, a managed set, or a storage system, you must fail over *all* components of the group, set, or system, respectively. Therefore, if only one component has failed, repairing that single component may be preferable to performing a complete failover.

The planned or unplanned failover of one or more DR groups should not be performed more frequently than once every 15 minutes. The planned or unplanned failover of a controller should also not be performed more frequently than once every 15 minutes.

Table 7 outlines example situations that call for a failover and those that do not. For each type of failover, a recommended action is described, which may require action at the source or destination storage system. Since replication can be bidirectional, one storage system can be the source and destination for different DR groups. You can use this table to customize contingency plans within your specific environment.

Table 7: When and when not to fail over a DR group, managed set, or storage system

Type of failure	Recommended action	
	DR group in normal mode	DR group in failsafe mode
Total loss of source storage system	Manual intervention to fail over data and processing to the destination storage system	Manual intervention to fail over data and processing to the destination storage system
Loss of both source controllers	Manual intervention to fail over data to destination storage system and then restart processing at the destination storage system	Manual intervention to fail over data to destination storage system and then restart processing at the destination storage system
Loss of single source controller	Failover not necessary	Failover not necessary
Total destination storage system loss	Failover not necessary	Manually intervene to continue processing at source storage system
Loss of both destination controllers	Failover not necessary	Manually intervene to continue processing at source storage system
Loss of all intersite links	Failover not necessary	Failover not necessary
Loss of both source intersite switches	Manual intervention to fail over data to destination storage system and then restart processing at both storage systems	Manual intervention to fail over data to destination storage system and then restart processing at both storage systems
Loss of single source intersite switch	Failover not necessary	Failover not necessary
Extended power outage at primary site	Manual intervention to fail over data and processing to remote site	Manual intervention to fail over data and processing to remote site
Loss of a Storage Management Appliance	Failover not necessary. Browse to standby appliance.	Failover not necessary. Browse to standby appliance.

Table 7: When and when not to fail over a DR group, managed set, or storage system (Continued)

Type of failure	Recommended action	
	DR group in normal mode	DR group in failsafe mode
Loss of single disk in redundant storage	Failover not necessary	Failover not necessary
Loss of single host of cluster	Failover not necessary	Failover not necessary
Disk group hardware failure (loss of redundancy) on the source storage system	Fail over to destination, and use procedure titled "Disk group hardware failure on the source storage system" on page 143	Fail over to destination, and use procedure titled "Disk group hardware failure on the source storage system" on page 143
Disk group hardware failure (loss of redundancy) on the destination storage system	Failover not necessary. Use procedure titled "Disk group hardware failure on the destination storage system" on page 150.	Failover not necessary. Use procedure titled "Disk group hardware failure on the destination storage system" on page 150.

The Continuous Access user interface

The Continuous Access user interface is a Java-based application that resides on the HP OpenView Storage Management Appliance (SMA) to simplify the performance of Continuous Access EVA procedures. It is discussed here for its role in performing failover-related procedures. Command View EVA also performs failover tasks. For information on using the Command View EVA application, refer to the HP StorageWorks Command View EVA Online Help.

Continuous Access user interface icons

All objects in the Continuous Access user interface showing storage systems and managed set areas are depicted by icons for an easy visual reference to the type of storage and its state. State is depicted by color, which indicates whether errors or warnings have been reported for that object. The meaning of the icon colors are:

- Green—The object and all objects in its tree report no errors.
- Yellow—The object or at least one object in its tree reports a warning.
- Red—The object or at least one object in its tree reports an error.
- White—The “health” of the object or one object in its tree is unknown.

[Table 8](#) describes the various icons encountered with the Continuous Access user interface. In the table, [status] is a placeholder for the status that the Continuous Access user interface reports when you rest your cursor over the icon. For example, the table may show “Copy Set: [status], Copying” and the result of placing the cursor on the icon may be “Copy Set: Normal, Copying.” If the Continuous Access user interface cannot determine the status of an object, it will report the status as an “unknown state.”

With a little experience, you can quickly identify the DR group icons without using the mouse cursor to initiate the ToolTip help. The following is guidance on understanding the DR group icons:

- The left box in the pair indicates the storage state being queried. For example, if you are working with storage system “HSV05,” the left box reflects its state, and the right box reflects the destination storage state.
- The direction of the arrow indicates which storage is acting as the source or destination. Remember that the source always replicates to the destination, so the arrow will point toward the destination.
- The Home indicator in a box indicates whether the storage is the “Home” or preferred storage system. If an arrow points to the Home system, then it is failed over.

- A log disk attached to a box indicates that either logging or merging is occurring. An arrow toward a log disk indicates logging. An arrow toward the storage and away from the log disk indicates merging.
- A simple broken arrow between two storage systems (boxes) indicates a failsafe-locked condition.
- A broken arrow between two storage systems (boxes) with logging shown could indicate an intersite link failure.

Table 8: Continuous Access user interface icons











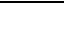






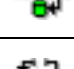







Icon	Description
Storage system and managed set icons	
	Storage System Group: <i>[status]</i>
	Storage System: <i>[status]</i>
	Managed Set: <i>[status]</i>
DR group icons	
	DR group: <i>[status]</i> , Source
	DR group: <i>[status]</i> , Destination
	DR group: <i>[status]</i> , Destination, Failed-over
	DR group: <i>[status]</i> , Source, Failed-over
	DR group: <i>[status]</i> , Source, Merging
	DR group: <i>[status]</i> , Destination, Merging
	DR group: <i>[status]</i> , Source, Failed-over, Merging
	DR group: <i>[status]</i> , Destination, Failed-over, Merging

Table 8: Continuous Access user interface icons (Continued)

Icon	Description
	DR group: <i>[status]</i> , Source, Suspended
	DR group: <i>[status]</i> , Destination, Suspended
	DR group: <i>[status]</i> , Destination, Failed-over, Suspended
	DR group: <i>[status]</i> , Source, Failed-over, Suspended
	DR group: <i>[status]</i> , Source, Logging—not-suspended (possible link failure)
	DR group: <i>[status]</i> , Destination, Logging—not-suspended (possible link failure)
	DR group: <i>[status]</i> , Destination, Failed-over, Logging—not-suspended (possible link failure)
	DR group: <i>[status]</i> , Source, Failed-over, Logging—not-suspended (possible link failure)
	DR group: <i>[status]</i> , Source, Failsafe-locked
	DR group: <i>[status]</i> , Destination, Failsafe-locked
	DR group: <i>[status]</i> , Destination, Failed-over, Failsafe-locked
	DR group: <i>[status]</i> , Source, Failed-over, Failsafe-locked
Copy set icons	
	Copy Set: <i>[status]</i>
	Copy Set: <i>[status]</i> , Copying

Data replication using the Continuous Access user interface

After the Continuous Access user interface is installed and initialized, active storage system objects are displayed in the left pane. Selecting an object and using the Data Replication menu (Figure 43) allows you to perform the following data replication functions:

- Suspend
- Resume
- Failover
- Disable Failsafe

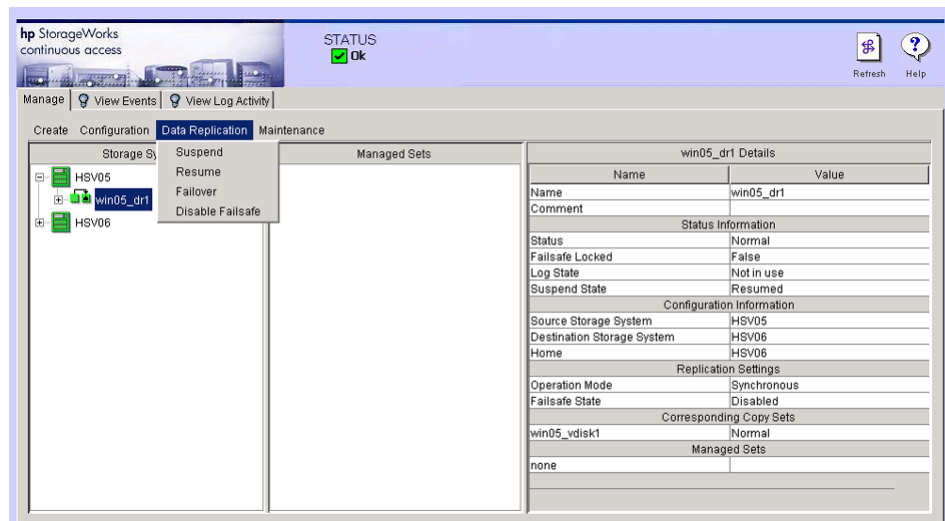


Figure 43: Data Replication menu on the Continuous Access user interface

The data replication function you select in the menu applies to the storage item selected, or to any item within its tree. For example, if failover is selected for a storage system, every DR group within the storage system performs a failover. If only a DR group within a storage system is selected for failover, only that DR group fails over.

Suspend

The Suspend menu item stops I/O between the source and destination DR group, managed set, or storage system. Logging begins when the Suspend command is issued. If failsafe is enabled on an object, then the Suspend command cannot be issued.

To suspend I/O for a DR group, managed set, or storage system:

1. Choose the object for which I/O is to be suspended.
2. Choose **Data Replication > Suspend** from the Continuous Access user interface main window menu bar (see Figure 43).
3. Verify that the object icon changes to indicate that I/O is suspended.

Resume

The Resume menu item restarts I/O between the source and destination DR group, managed set, or storage system. A merge of the log is also performed to synchronize the data between the source and destination copy sets. If the log is full, a full copy is performed.

To resume I/O for a DR group, managed set, or storage system:

1. Choose the object for which I/O is to be resumed.
2. Choose **Data Replication > Resume** from the Continuous Access user interface main window menu bar (see Figure 43).
3. Verify that the object icon changes to indicate that I/O is resumed.

Failover

The Failover menu item reverses the source and destination roles for the DR group, managed set, or storage system. Although transparent to the user, the command is issued to the destination storage system. The host presentations for the storage system becoming the new source system are immediately enabled for read-write access with the same access rights used by the previous source storage system.

To fail over a DR group, managed set, or storage system:

1. Choose the object for which failover is desired.
2. Choose **Data Replication > Failover** from the Continuous Access user interface main window menu bar (see Figure 43).
3. Verify that the object icon changes to indicate that failover has occurred.

Disable Failsafe

The Disable Failsafe menu item enables a DR group, managed set, or storage system to allow logging if the destination becomes inaccessible.

To disable failsafe for a DR group, managed set, or storage system:

1. Choose the object for which the failsafe mode is to be disabled.
2. Choose **Data Replication > Disable Failsafe** from the Continuous Access user interface main window menu bar (see Figure 43).
3. Verify that the object icon changes to indicate that it is not failsafe-locked.

Possible event scenarios

When a situation occurs that requires manual intervention in the data replication process, one of the following scenarios will in all likelihood describe your current situation or your desired action:

- [Planned failover](#)
- [Unplanned failover](#)
- [Resumption of operations if unable to access destination while source in failsafe-locked state \(extended period of time\)](#)
- [Return operations to home storage system](#)
- [Return operations to replaced new storage hardware](#)
- [Disk group hardware failure on the source storage system](#)
- [Disk group hardware failure on the destination storage system](#)

The procedures for performing these actions are described later in this chapter.

Planned failover

Situation: Due to scheduled maintenance at the primary site, you need to perform a planned move of operations from the source storage system to the destination storage system.

Action: Prepare the source storage system for the failover, and then perform a failover to the destination storage system. After the failover is complete, you can continue to operate from this storage system and revert back to failsafe mode, if desired. When the maintenance is complete you can failover to the original source storage system (see the procedure “Planned failover” on page 119.)

Note: *Home* is a designation you can set in the Continuous Access user interface to identify the preferred storage system. By default, a storage system created as a source is designated as the Home storage system. Because the role of the source storage system can change during failover, this Home designation allows you to identify your preferred storage system if you choose to do so.

Unplanned failover

Situation: You have experienced an unplanned loss of the primary site or the source storage system. The duration of the outage at the source is unknown. The Continuous Access EVA hardware components (hosts, controllers, and switches, for example) at the primary site may or may not remain intact.

Action: Perform an immediate failover to the remote site or passive storage system. When the primary site is back online, you can choose to return to the Home storage system or to one with new hardware (see the procedure “Unplanned failover” on page 127.)

Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)

Situation: You have experienced an unplanned loss of the destination storage system, or the connection to the destination storage system, due to failure of the intersite links, loss of power at the alternate site, loss of both destination switches, and so on. The duration of the outage is unknown. The DR groups are in failsafe mode and host I/O is paused because the DR groups are failsafe-locked.

Action: Change from failsafe to normal mode, and then resume host I/O until the connection to the destination storage system is re-established. When the connection to the destination site is stable, perform a change back to the failsafe mode (see the procedure “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132).

Return operations to home storage system

Situation: You are operating from a storage system that was not originally designated as Home within the Continuous Access user interface. You need to perform a planned move of operations from this alternate source storage system to the Home storage system.

Action: Prepare the Home storage system for the failover, and then perform a failover to the Home storage system (see the procedure “Return operations to home storage system” on page 134).

Return operations to replaced new storage hardware

Situation: Some type of disaster (lightning, flood, fire, severe equipment failure, or the like) has damaged equipment at the Home site and forced a failover to an alternate site. You are operating from a storage system that was not originally designated as Home within the Continuous Access user interface.

Action: When the damaged components at the Home site (hosts, controllers, or switches, for example) have been repaired, and the site is operational and back online, perform a failover to new hardware at the Home site (see the procedure “Return operations to replaced new storage hardware” on page 134).

Disk group hardware failure on the source storage system

Situation: A hardware failure on your source storage system causes a disk group to become inoperative. This can be caused by the loss of enough disks to create a loss of redundancy within the disk group and affect all Vraid types present on the disk group.

Action: If you plan to recover using data on the destination storage system, then failover to the destination storage system. Delete DR groups and Vdisks on the failed storage system. Repair the failed disk group. Re-create DR groups, Vdisks, and host presentations.

If the failed source storage system was logging at the time of the hardware failure, then you must choose to recover with data at the destination site or from a backup.

Disk group hardware failure on the destination storage system

Situation: A hardware failure on your destination storage system causes a disk group to become inoperative. This can be caused by the loss of enough disks to create a loss of redundancy within the disk group and affect all Vraid types present on the disk group.

Action: Delete the DR groups on the source storage system that replicated to the failed disk group. Repair the failed disk group on the destination storage system. Re-create your DR groups on the source storage system and make host presentations at the destination storage system.

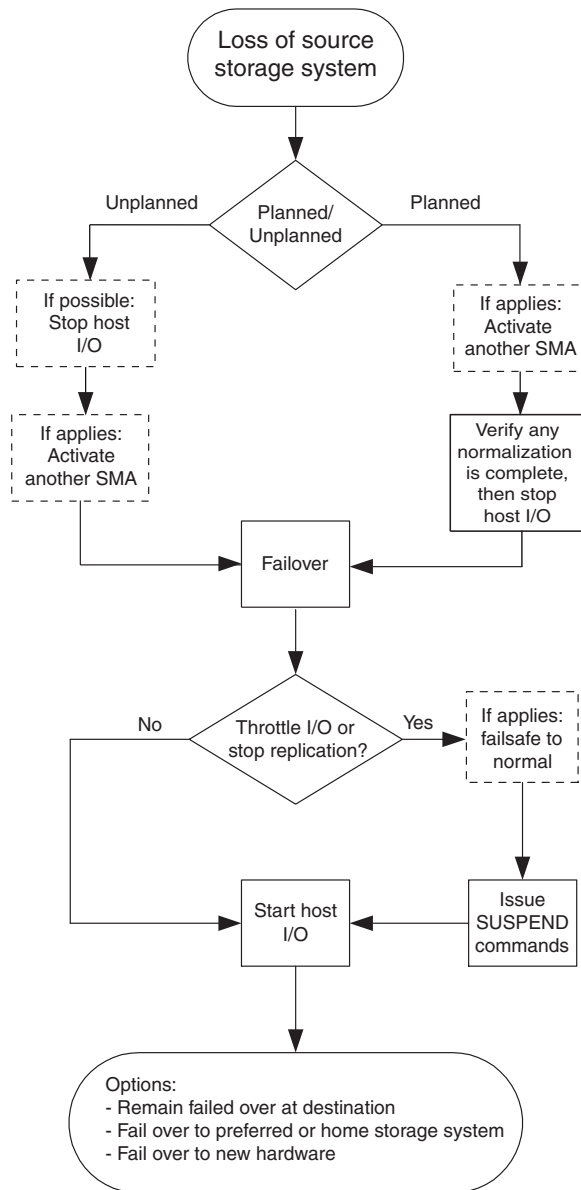
Failover and recovery procedures

The following detailed procedures describe the steps required to resolve the scenarios mentioned above. HP recommends that you rehearse or practice the procedures in this chapter so that you will be prepared to perform a failover, resumption of operation, or hardware recovery procedure quickly and accurately during a crisis. You may want to customize these procedures for your own use.

Planned failover

See [Figure 44](#) for the flow of steps required for a planned transfer of operations to a remote site. The following procedure describes these steps in more detail:

1. If desired, move storage management to another SMA.
2. Check to ensure that full normalization has occurred. If a merge or full copy was occurring when the failover process started, wait for them to complete.
3. Stop all host I/O on the source storage system. Follow the steps listed below for each operating system in your heterogeneous configuration:
 - a. **HP OpenVMS:** If the operating system is up and running, and is being used exclusively for Continuous Access EVA operations, shut down the operating system and power off the hosts. If the operating system is being used for other applications, remove all I/O to the Vdisks that will be failed over, and then unmount the volumes associated with these Vdisks.
 - b. **HP Tru64 UNIX:** If the operating system is up and running and is being used exclusively for Continuous Access EVA operations, shut down the operating system and power off the hosts. If the operating system is being used for other applications, remove all I/O and unmount all file system Vdisks that will be failed over.
 - c. **HP-UX:** If the operating system is up and running, remove all I/O to the Vdisks that will be failed over, and then unmount the file systems associated with the Vdisks.
 - d. **IBM AIX:** If the operating system is up and running, remove all I/O to the Vdisks that will be failed over, then unmount the file systems associated with the Vdisks.



CXO8065B

Figure 44: Planned and unplanned transfer of operations

- e. **Linux:** If the operating system is up and running, remove all I/O to the Vdisks that will be failed over, and then unmount the file systems associated with the Vdisks.

If you are running Logical Volume Manager (LVM) with or without clustering, see “Bootless DR group planned failover with Linux using LVM in standalone mode or with SuSE SLES 8 running LifeKeeper 4.4.3” on page 197.
 - f. **Microsoft Windows NT-X86:** If the operating system is up and running, shut it down.
 - g. **Microsoft Windows 2000/2003:** If the operating system is up and running, shut it down.
 - h. **Novell NetWare:** If the operating system is up and running, remove all I/O to the Vdisks that will be failed over, and then dismount the volumes associated with these Vdisks.
 - i. **Sun Solaris:** If the operating system is up and running and is being used exclusively for Continuous Access EVA operations, shut down the operating system and power off the hosts. If the operating system is being used for other applications, remove all I/O and unmount all volumes that have Vdisks that will be failed over.
- 4. Perform a failover to the alternate site.
 - 5. If you plan to throttle I/O to specific storage systems, suspend your less important DR groups at your new source. This will force the controllers to replicate the most important data first when the links to the previous source controller are re-established.
 - 6. If you plan to operate for an extended time at the alternate site (Home storage system and Fibre Channel links must be functioning properly) and you have a DR group that needs failsafe mode enabled, perform these steps:
 - a. If DR groups were suspended, resume copying on affected destination DR groups. Wait for the log disk to finish merging.
 - b. Change affected DR groups to failsafe mode.

Note: You can enable failsafe mode at the destination storage system while a merge or full copy is being performed.

7. Issue operating system–dependent commands for presentation of units to remote hosts to start host I/O.
 - a. **HP OpenVMS:** Allow hosts to recognize new units:
 - 1) If the remote hosts are shut down, boot them now. Booting the hosts enables OpenVMS to recognize the drives.
 - 2) If the remote hosts are not shut down, use the following command from a privileged account to enable OpenVMS to recognize the drives:

```
MCR SYSMAN IO AUTOCONFIGURE/LOG
```

- b. **HP Tru64 UNIX:** Allow the hosts to recognize new units.
 - 1) If the remote hosts are shut down, boot them now. Booting the hosts enables Tru64 UNIX to recognize the drives.
 - 2) If the remote hosts are not shut down, use the following command to recognize the drives:

```
hwmgr -scan scsi
```

This may take awhile for large configurations. If this is the case, scan only the SCSI buses that have new units added. Scan only one bus at a time. Use the following command:

```
hwmgr -scan scsi -bus x
```

(where *x* is the SCSI bus number)

- c. **HP-UX:** Allow the remote hosts to recognize the new units.
 - 1) If the remote hosts are shut down, boot them now. Booting the hosts enables HP-UX to recognize the drives.
 - 2) If the remote hosts are not shut down, use the following command to enable HP-UX to recognize the drives and verify that they are present. This command will display only the previous configured failed-over LUNs:

```
ioscan -fnCdisk
```

If the device special files are not displayed, run `insf -e`, and then run `ioscan -fnCdisk` again.

Run the command:

```
vgimport VolumeGroupName DeviceSpecialFile
```

Repeat the previous command for each new failed-over LUN.

Use the following command to mount the LUNs:

```
mount -a
```

Note: *VolumeGroupName* is the name of the volume group you originally created at the local site. The *DeviceSpecialFiles* are from the ioscan in the form of `/dev/dsk/c_t_d/`.

For consistency, configure the same *DeviceSpecialFiles* with the same volume groups, logical volumes, and file systems for the failed-over LUNs at the remote site with the same LUNs at the local site.

d. **IBM AIX:** Allow the hosts to recognize new units.

- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables AIX to recognize the drives.
- 2) If the remote hosts are not shut down, use the following commands to recognize the drives and verify that they are present:

```
cfgmgr -v
lsdev -Cc disk
```

Use the following commands to access file systems on the failed-over Vdisks:

```
importvg -y VolumeGroupName hdiskx
mount all
```

Note: *VolumeGroupName* is the name of the volume group you originally created at the local site, and *x* is the number of the hdisk assigned to the failed-over Vdisk. If the `-y VolumeGroupName` parameter is omitted, AIX will create a default volume group name for you (for example, `vg00`).

- e. **Linux:** Reboot the servers at the remote site and then remount the file system.
- f. **Microsoft Windows NT:** Allow the remote hosts to recognize new units.
Reboot the servers at the remote site and log in using an account that has administrative privileges. You should be able to see all of the units in My Computer.

- g. **Microsoft Windows 2000/2003:** Allow the remote hosts to recognize new units.
 - 1) On each host, log in using an account that has administrative privileges.
 - 2) Open Computer Management and click Disk Management.
 - 3) After Disk Management has initialized, go to the Action Menu and click Rescan Disks. If the units fail to appear, click **F5** (Refresh). All of the failed-over units should appear in the right-hand pane.
- h. **Novell NetWare:** Allow the hosts to recognize new units.
 - 1) If the remote hosts are shut down, boot them now. If you are using traditional NetWare volumes, booting the hosts allows Novell NetWare to recognize the drives and automatically mount the volumes. If you are using NSS logical volumes, booting the hosts will recognize the NSS pools and activate them. However, you must manually mount each individual NSS volume by typing `MOUNT VolumeName` at the NetWare console.
 - 2) If the remote hosts are already up and running, or if they do not recognize the drives, issue the following command from the console before mounting the volumes:

```
SCAN FOR NEW DEVICES
```

Alternatively, you can use the *NWCONFIG* utility to issue this same command.

Next, mount the volumes with these commands:

```
MOUNT ALL (for traditional NetWare volumes)
```

```
MOUNT VolumeName (for NSS logical volumes)
```

- i. **Sun Solaris:** Allow the remote hosts to recognize new units.
 - 1) Reboot the remote hosts using the `reboot -- -r` command, or use the following version-dependent commands to update the Secure Path Manager:

Solaris 6, 7, and 8 (run the following commands):

```
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```


Solaris 9:

- a) Present new units with LUN numbers sequentially following the old LUNs.
- b) Run the following commands:

```
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

You should be able to see all of the units with two paths in the Secure Path Manager. You should also be able to see all of the units by using the `format` command.

- 2) If Secure Path was not configured for these units, use the following version-dependent commands to add them to the Secure Path Manager.

Solaris 6, 7, and 8 (run the following commands):

```
/opt/CPQswsp/bin/spconfig  
/opt/CPQswsp/bin/spmgr/display -u  
/opt/CPQswsp/bin/spmgr add WWLUNID  
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

Solaris 9:

- a) Add the units with `spmgr add WWLUNID` or `spmgr add-r WWNN all`.
- b) Run `update_drv -f sd` to inform the system about attribute changes to the `sd` driver.
- c) Run `disks` to create `/dev` entries for the new units.

You should now be able to see the drives using the `format` command. Refer to the current version of Secure Path documentation for additional assistance.

After the transfer of operation is complete, you have three options after the cause of the failover is resolved:

- Remain failed over at the alternate (or destination) site
- Return operations to the Home storage system (see procedure below)
- Return operations to new hardware (see procedure below)

Unplanned failover

See [Figure 44](#) for the flow of steps required for an unplanned transfer of operations to a remote site. The following procedure describes these steps in more detail:

1. If your hosts are running on the source storage system, and you are able to access these hosts, then stop all host I/O.
2. If you cannot access the SMA managing the storage systems, establish management control with another SMA. Refer to Chapter 5 for SMA procedures.
3. Perform a failover to the destination site.
4. If you plan to throttle I/O to specific storage systems, suspend your less important DR groups at your new source. This forces the controllers to replicate the most important data first when the links to the previous source controller are re-established.
5. Issue operating system dependent commands for presentation of units to remote hosts to start host I/O.

a. **HP OpenVMS:** Allow hosts to recognize new units:

- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables OpenVMS to recognize the drives.
- 2) If the remote hosts are not shut down, use the following command from a privileged account to enable OpenVMS to recognize the drives:

```
MCR SYSMAN IO AUTOCONFIGURE/LOG
```

b. **HP Tru64 UNIX:** Allow the hosts to recognize new units.

- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables Tru64 UNIX to recognize the drives.
- 2) If the remote hosts are not shut down, use the following command to recognize the drives:

```
hwmgr -scan scsi
```

This may take awhile for large configurations. If this is the case, scan only the SCSI buses that have new units added. Scan only one bus at a time. Use the following command:

```
hwmgr -scan scsi -bus x
```

(where *x* is the SCSI bus number)

c. **HP-UX:** Allow the remote hosts to recognize the new units.

- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables HP-UX to recognize the drives.
- 2) If the remote hosts are not shut down, use the following command to enable HP-UX to recognize the drives and verify that they are present. This command will display only the previous configured failed-over LUNs:

```
ioscan -fnCdisk
```

If the device special files were not displayed, run `insf -e`, and then run `ioscan -fnCdisk` again.

Run the command:

```
vgimport VolumeGroupName DeviceSpecialFile
```

Repeat the previous command for each new failed-over LUN.

Use the following command to mount the LUNs:

```
mount -a
```

Note: *VolumeGroupName* is the name of the volume group you originally created at the local site. The *DeviceSpecialFiles* are from the `ioscan` in the form of `/dev/dsk/c_t_d/`.

For consistency, configure the same *DeviceSpecialFiles* with the same volume groups, logical volumes, and file systems for the failed-over LUNs at the remote site with the same LUNs at the local site.

d. **IBM AIX:** Allow the hosts to recognize new units.

- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables AIX to recognize the drives.
- 2) If the remote hosts are not shut down, use the following commands to recognize the drives and verify that they are present:

```
cfgmgr -v
```

```
lsdev -Cc disk
```

Use the following commands to access file systems on the failed-over Vdisks:

```
importvg -y VolumeGroupName hdiskx  
mount all
```

Note: *VolumeGroupName* is the name of the volume group you originally created at the local site, and x is the number of the hdisk assigned to the failed-over Vdisk. If the -y VolumeGroupName parameter is omitted, AIX will create a default volume group name for you, for example, vg00.

- e. **Linux:** Reboot the servers at the remote site and then remount the file system.
- f. **Microsoft Windows NT:** Allow the remote hosts to recognize new units.
Reboot the servers at the remote site and log in using an account that has administrative privileges. You should be able to see all of the units in My Computer.
- g. **Microsoft Windows 2000/2003:** Allow the remote hosts to recognize new units.
 - 1) On each host, log in using an account that has administrative privileges.
 - 2) Open Computer Management and click Disk Management.
 - 3) After Disk Management has initialized, go to the Action Menu and click Rescan Disks. If the units fail to appear, click **F5** (Refresh). All of the failed-over units should appear in the right-hand pane.
- h. **Novell NetWare:** Allow the hosts to recognize new units.
 - 1) If the remote hosts are shut down, boot them now. If you are using traditional NetWare volumes, booting the hosts allows Novell NetWare to recognize the drives and automatically mount the volumes. If you are using NSS logical volumes, booting the hosts will recognize the NSS pools and activate them. However, you must manually mount each individual NSS volume by typing `MOUNT VolumeName` at the NetWare console.

- 2) If the remote hosts are already up and running, or if they do not recognize the drives, issue the following command from the console before mounting the volumes:

```
SCAN FOR NEW DEVICES
```

Alternatively, you can use the *NWCONFIG* utility to issue this same command.

Next, mount the volumes with these commands:

```
MOUNT ALL (for traditional NetWare volumes)
```

```
MOUNT VolumeName (for NSS logical volumes)
```

- i. **Sun Solaris:** Allow the remote hosts to recognize new units.

- 1) Reboot the remote hosts using the `reboot -- -r` command, or use the following version-dependent commands to update the Secure Path Manager:

Solaris 6, 7, and 8 (run the following commands):

```
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

Solaris 9:

- a) Present new units with LUN numbers sequentially following the old LUNs.
- b) Run `drvconfig:disks` to be able to see the devices in the `spmgr display -u list`.

You should be able to see all of the units with two paths in the Secure Path Manager. You should also be able to see all of the units by using the `format` command.

- 2) If Secure Path was not configured for these units, use the following version-dependent commands to add them to the Secure Path Manager.

Solaris 6, 7, and 8 (run the following commands):

```
/opt/CPQswsp/bin/spconfig  
/opt/CPQswsp/bin/spmgr/display -u  
/opt/CPQswsp/bin/spmgr add WWLUNID  
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

Solaris 9:

- a) Add the units with `spmgr add WWLUNID` or `spmgr add-r WWNN all`.
- b) Run `update_drv -f sd` to inform the system about attribute changes to the sd driver.
- c) Run `disks` to create `/dev` entries for the new units.

You should now be able to see the drives using the `format` command. Refer to the current version of Secure Path documentation for additional assistance.

With Solaris, you may need to execute the `fsck` command before the operating system can mount the LUN. HP recommends that you run the `fsck` command with the `-m` option before running `fsck` to repair the file system.

When the transfer of operation is complete and the cause of the failover is resolved, you have three options:

- Remain failed over at the alternate (or destination) site
- Return operations to the Home storage system (see procedure below)
- Return operations to new hardware (see procedure below)

Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)

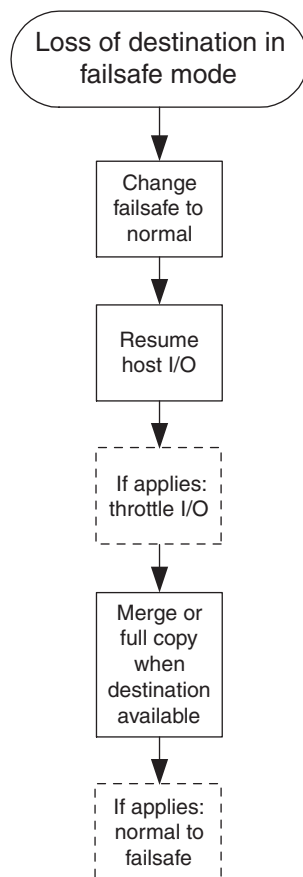
See [Figure 45](#) for the flow of steps required to resume operations if you are unable to access the destination while in a failsafe-locked state. The following procedure describes these steps in more detail:

1. Change affected source DR groups from failsafe mode to normal mode.
2. If necessary, issue operating system dependent commands to the local hosts to start I/O again on the units that were failsafe-locked.
3. If you plan to throttle I/O to specific storage systems, suspend your less important DR groups. This forces the controllers to replicate the most important data first when the links are re-established. When ready to merge to destination from source, issue the Resume command.

Note: If you stay in suspended mode for an extended length of time, you run the risk of overrunning the log. A log overrun initiates a full copy. During a full copy, the data is not usable until the full copy completes.

4. When connections to the destination site are re-established and merging is complete, change DR groups from normal mode to failsafe mode, if desired.

Note: If source DR groups go into full copy mode, you can also enable failsafe mode.



CXO8066A

Figure 45: Resumption of operations if unable to access destination in failsafe mode

Return operations to home storage system

The return operations back to the home storage system procedure is similar to a planned failover. Perform the following steps:

1. From the Continuous Access user interface, click the **Refresh** icon and select **Discovery**.
2. Perform the procedure “Planned failover” on page 119.

Return operations to replaced new storage hardware

This procedure is used after a failure that resulted in the replacement of storage system hardware at what had been the source storage system. The procedure does not include steps to rebuild servers using the storage (this should be part of your overall disaster plan). The new hardware is now acting as the destination storage system after a failover and is referred to in this procedure as the system with new hardware, or the storage system with failed hardware. The surviving system is now your source storage system after the failover. [Figure 46](#) and the steps below explain the process required to return operations to a system having replaced new hardware:

1. Denote your storage system names having failed or new hardware (destination) and your surviving storage system (source) in the table provided below. For example, your storage system with new hardware may be named HSV01 and your surviving storage system may be named HSV02. The table can be used as a reference while continuing with the procedural steps.

	Storage system with failed or new hardware	Surviving storage system
Storage System Name		
Storage System Name		
Storage System Name		
Storage System Name		
Storage System Name		
Storage System Name		
Storage System Name		
Storage System Name		

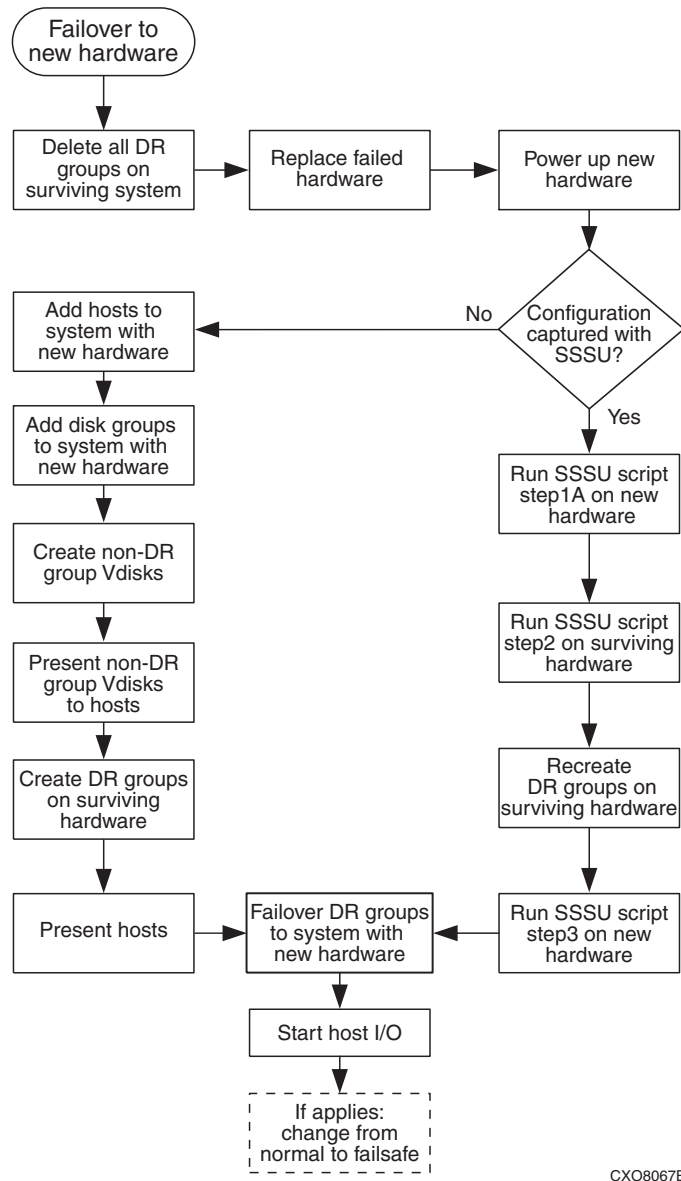


Figure 46: Return operations to new hardware

2. Using the Continuous Access user interface, click the **Refresh** icon. A Discover window opens. Select **Refresh Display**. A refresh may take several minutes depending on the size of the configuration and the distance between storage systems.
3. Delete all DR groups on the surviving system that ever had a relationship with the failed hardware.
4. Replace the failed hardware. Depending on the failure, this means replacing hard drives or controllers, deleting disk groups, and so on.
5. Remove the connection between the source and destination storage systems. This can be accomplished by removing it from the SAN, disabling the intersite links, or by placing the storage systems into separate zones.

To place the storage system into separate zones, you need two zones. One zone contains the source storage system, source hosts, and the SMA. The second zone contains the destination storage systems, destination hosts, and the SMA.
6. If there are any destination DR groups on the previously failed storage system, delete them with Command View EVA. Click the Data Replication folder, select a DR group, and click **Delete**. If this is not successful, the source-destination connection still exists, so go to the previous step.
7. Delete all Vdisks that were members of DR groups on the destination storage system.
8. Re-establish communication between the source and destination storage systems. Either add the storage system back into the SAN, enable the intersite links, or place the storage systems into the same zone.
9. Rebuild your configurations of Vdisks and DR groups.
10. Rediscover your storage with the Continuous Access user interface.
11. (This step is optional.) To preserve your existing zoning, give the new hardware the same World Wide Names as they existed with the failed hardware.

12. Perform one of the following:
 - a. If the replaced storage system configuration was captured with the Storage System Scripting Utility (SSSU), then execute the script *ConfigName_step1A* on the new hardware, and proceed to step 13. *ConfigName* is a user-assigned name given to the SSSU script at the time of creation. (See the procedure titled “Backing Up Configuration Information” on page 82.)
 - b. If you are not using an SSSU script for recovery, initialize the newly replaced storage system.
13. Add in your hosts for the system with new hardware.
14. Add in your disk groups on the new hardware.
15. Create your non-DR group Vdisks.
16. Present all non-DR group Vdisks to their hosts.
17. Perform one of the following:
 - a. If the surviving storage system configuration was captured with the SSSU, then execute *ConfigName_step2* on the surviving storage system. *ConfigName* is a user-assigned name given to the SSSU script at the time of creation. DR groups are re-created with the SSSU if they were performing as the source when the configuration was captured. This step may take some time to complete.
 - b. If you are not using an SSSU script for recovery, re-create all DR groups on the surviving system with destinations set to the new system.
18. If you used the SSSU to re-create DR groups on the surviving storage system in the previous step, then you must manually re-create any additional DR groups on the surviving storage system that had their source on the failed hardware. This is necessary because the SSSU will not re-create those DR groups on the surviving storage system if they performed as the destination when the configuration was captured. After you perform this step, all DR groups now reside on the surviving storage system.
19. Perform one of the following:
 - a. If the original storage system configuration was captured with the SSSU, then execute *ConfigName_step3* on the new hardware. *ConfigName* is a user-assigned name given to the SSSU script at the time of creation.
 - b. If you are not using an SSSU script for recovery, present the destination Vdisks on the system with new hardware to the appropriate hosts.

20. If you used the SSSU to present destination Vdisks to their hosts in the previous step, you must manually present any additional Vdisks to their hosts on the storage system with new hardware that originally had their sources on the failed hardware.

The reason for this manual presentation is because the SSSU will not present Vdisks whose destinations were to the surviving storage system when the configuration was captured. After performing this step, all destination Vdisks have been presented to hosts.

21. Fail over any DR groups to the new storage system using the procedure “Planned failover” on page 116 if this system is to be the Home system, or if the system is to be the source for the DR groups.

22. Issue operating system–dependent commands for presentation of units to hosts to start I/O:

- a. **HP OpenVMS:** Allow hosts to recognize new units:

- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables OpenVMS to recognize the drives.
- 2) If the remote hosts are not shut down, use the following command from a privileged account to enable OpenVMS to recognize the drives:

```
MCR SYSMAN IO AUTOCONFIGURE/LOG
```

- b. **HP Tru64 UNIX:** Allow the hosts to recognize new units.

- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables Tru64 UNIX to recognize the drives.
- 2) If the remote hosts are not shut down, use the following command to recognize the drives:

```
hwmgr -scan scsi
```

This may take awhile for large configurations. If this is the case, scan only the SCSI buses that have new units added. Scan only one bus at a time. Use the following command:

```
hwmgr -scan scsi -bus x
```

(where x is the SCSI bus number)

- c. **HP-UX:** Allow the remote hosts to recognize the new units.
- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables HP-UX to recognize the drives.
 - 2) If the remote hosts are not shut down, use the following command to enable HP-UX to recognize the drives and verify that they are present. This command will display only the previous configured failed-over LUNs:

```
ioscan -fnCdisk
```

If the device special files were not displayed, run `insf -e`, then run `ioscan -fnCdisk` again.

Run the command:

```
vgimport VolumeGroupName DeviceSpecialFile
```

Repeat the previous command for each new failed-over LUN.

Use the following command to mount the LUNs:

```
mount -a
```

Note: *VolumeGroupName* is the name of the volume group you originally created at the local site. The *DeviceSpecialFiles* are from the `ioscan` in the form of `/dev/dsk/c_t_d/`.

For consistency, configure the same *DeviceSpecialFiles* with the same volume groups, logical volumes, and file systems for the failed-over LUNs at the remote site with the same LUNs at the local site.

- d. **IBM AIX:** Allow the hosts to recognize new units.
- 1) If the remote hosts are shut down, boot them now. Booting the hosts enables AIX to recognize the drives.
 - 2) If the remote hosts are not shut down, use the following commands to recognize the drives and verify that they are present:

```
cfgmgr -v
```

```
lsdev -Cc disk
```

Use the following commands to access file systems on the failed-over Vdisks:

```
importvg -y VolumeGroupName hdiskx  
mount all
```

Note: *VolumeGroupName* is the name of the volume group you originally created at the local site, and *x* is the number of the hdisk assigned to the failed-over Vdisk. If the *-y VolumeGroupName* parameter is omitted, AIX will create a default volume group name for you, for example, vg00.

- e. **Linux:** Reboot the servers at the remote site and then remount the file system.
- f. **Microsoft Windows NT:** Allow the remote hosts to recognize new units.
Reboot the servers at the remote site and log in using an account that has administrative privileges. You should be able to see all of the units in My Computer.
- g. **Microsoft Windows 2000/2003:** Allow the remote hosts to recognize new units.
 - 1) On each host, log in using an account that has administrative privileges.
 - 2) Open Computer Management and click Disk Management.
 - 3) After Disk Management has initialized, go to the Action Menu and click Rescan Disks. If the units fail to appear, click **F5** (Refresh). All of the failed-over units should appear in the right-hand pane.
- h. **Novell NetWare:** Allow the hosts to recognize new units.
 - 1) If the remote hosts are shut down, boot them now. If you are using traditional NetWare volumes, booting the hosts allows Novell NetWare to recognize the drives and automatically mount the volumes. If you are using NSS logical volumes, booting the hosts will recognize the NSS pools and activate them. However, you must manually mount each individual NSS volume by typing `MOUNT VolumeName` at the NetWare console.

- 2) If the remote hosts are already up and running, or if they do not recognize the drives, issue the following command from the console before mounting the volumes:

```
SCAN FOR NEW DEVICES
```

Alternatively, you can use the *NWCONFIG* utility to issue this same command.

Next, mount the volumes with these commands:

```
MOUNT ALL (for traditional NetWare volumes)
```

```
MOUNT VolumeName (for NSS logical volumes)
```

- i. **Sun Solaris:** Allow the remote hosts to recognize new units.

- 1) Reboot the remote hosts using the `reboot -- -r` command, or use the following version-dependent commands to update the Secure Path Manager:

Solaris 6, 7, and 8 (run the following commands):

```
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

Solaris 9:

- a) Present new units with LUN numbers sequentially following the old LUNs.
- b) Run `drvconfig:disks` to be able to see the devices in the `spmgr display -u list`.

You should be able to see all of the units with two paths in the Secure Path Manager. You should also be able to see all of the units by using the `format` command.

- 2) If Secure Path was not configured for these units, use the following version-dependent commands to add them to the Secure Path Manager.

Solaris 6, 7, and 8 (run the following commands):

```
/opt/CPQswsp/bin/spconfig  
/opt/CPQswsp/bin/spmgr/display -u  
/opt/CPQswsp/bin/spmgr add WWLUNID  
drvconfig -v  
disks  
/opt/CPQswsp/bin/spmgr display
```

Solaris 9:

- a) Add the units with `spmgr add WWLUNID` or `spmgr add-r WWNN all`.
- b) Run `update_drv -f sd` to inform the system about attribute changes to the sd driver.
- c) Run `disks` to create `/dev` entries for the new units.

You should now be able to see the drives using the `format` command. Refer to the current version of Secure Path documentation for additional assistance.

23. At this point, you have the option of setting all affected DR groups from normal mode to failsafe mode.

Recovering from a disk group hardware failure

The condition referred to as a *disk group hardware failure* occurs when a disk group loses a quantity of disks beyond the capability from which a given Vraid type can recover. It is a loss of redundancy that results in an inoperative disk group. This condition can occur from the loss of one disk for Vraid0 to as few as two disks for Vraid1 and Vraid5. In each case, the hardware failure needs to be fixed, and the disk group data has to be structurally rebuilt. This section describes the symptoms and recovery methods for an inoperative disk group at either the source and destination storage systems.

Disk group hardware failure on the source storage system

A disk group hardware failure on the source storage system can result in two different recovery methods:

- If data replication was occurring normally when the source disk group became inoperative, then the data at the destination storage system is current. A failover is performed to the destination storage system, DR groups are deleted, the inoperative disk group is repaired, and the DR groups are re-created. Data is then copied back.
- If your disk group becomes inoperative when your DR groups are logging (for example, your DR groups were suspended, or the intersite links are down), then your data is *stale* on the destination storage system. Stale data is older data that is not as current as what exists on its partnered storage system. If you prefer to use stale data for recovery, then the steps are the same as if replication was occurring normally. However, if you prefer to continue from a point-in-time, then the inoperative disk group is repaired, and data is restored from a backup or full copy.

Note: When you delete DR groups to recover from a disk group hardware failure, you lose the redundancy of the other site or disaster tolerance of your data.

Recovery when data replication was normal before failure

This procedure is performed when a disk group hardware failure occurs on the source storage system and the data on the destination storage system is current. Other than some initial failure indications you may see with the Continuous Access user interface, most of these recovery procedures are performed with Command View EVA.

[Figure 47](#) shows a healthy storage system when viewed by the Continuous Access user interface. All storage is green and the overall status is OK.

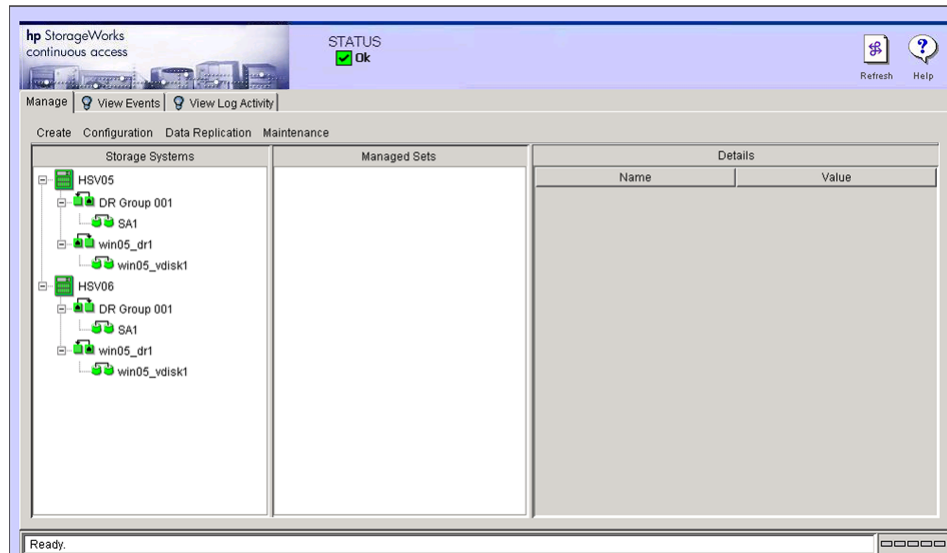


Figure 47: Normal disk group indication with Continuous Access user interface

When a disk group becomes inoperative, the Continuous Access user interface status changes to unknown and your storage in that disk group becomes unavailable. [Figure 48](#) shows that a source DR group named win05_dr1 that is in storage system HSV05 has become inoperative. The page status now is **Unknown**, and the value of the corresponding DR group shows **Unknown**.

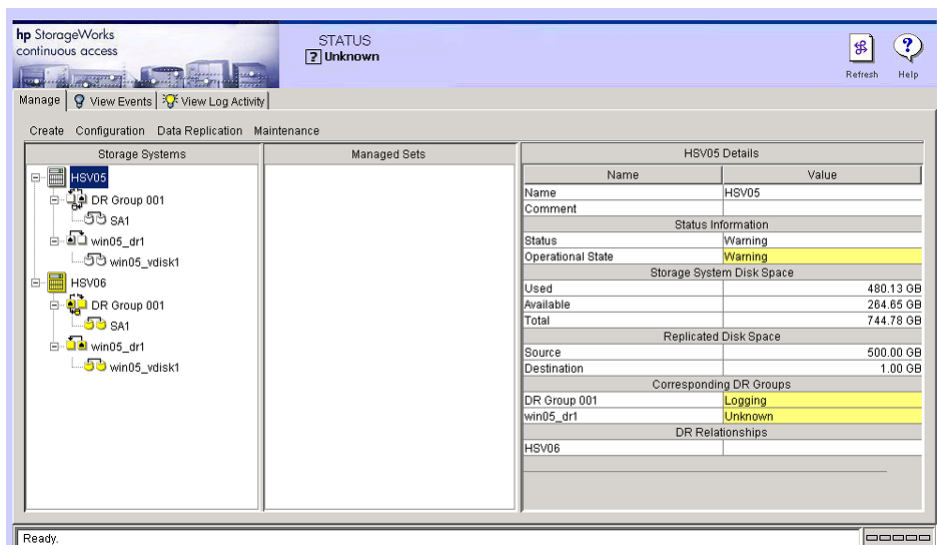


Figure 48: Hardware failure viewed from Continuous Access user interface

Figure 49 shows a failed disk group when viewed by Command View EVA.

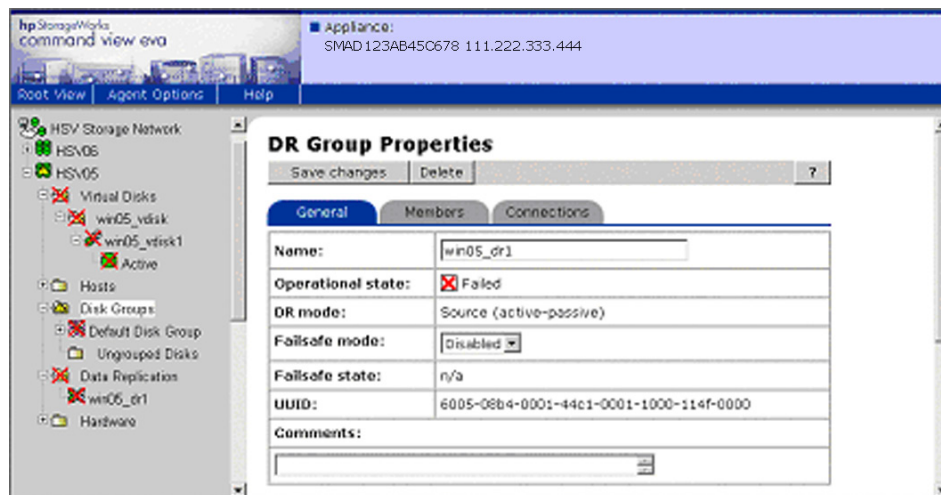


Figure 49: Hardware failure viewed from Command View EVA

When you receive an indication that a disk group on your source storage system is inoperative, perform the following steps using Command View EVA:

1. Navigate to each DR group on the surviving storage system and fail them over.
2. Navigate to each DR group on your surviving storage system that was failed over and delete each DR group. DR groups will still remain on the failed storage system, but will be deleted in later steps.
3. Navigate to the failed disk group. A screen similar to [Figure 50](#) is displayed and lists failed Vdisks and DR groups. Click the **Start deletion process** tab.

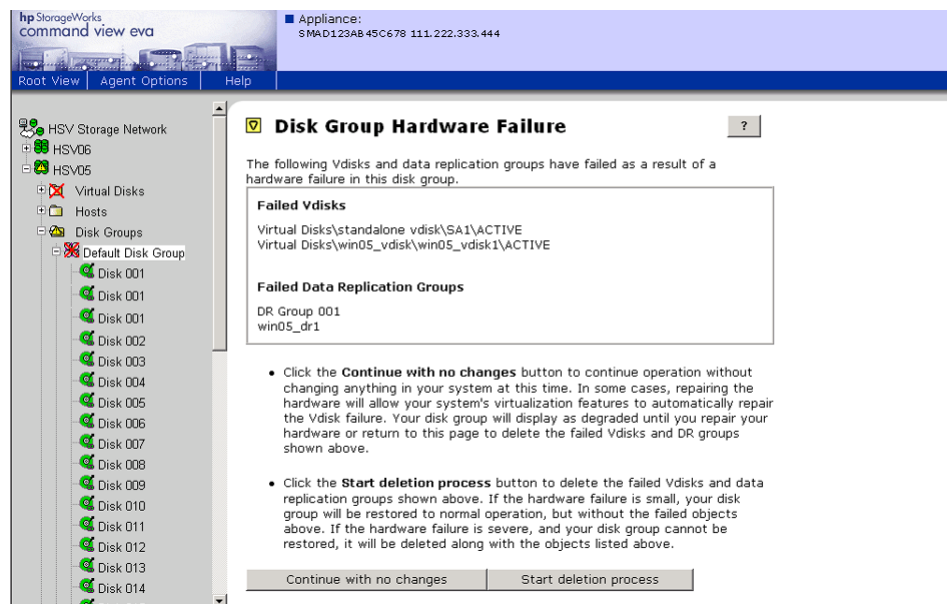


Figure 50: Command View EVA hardware deletion process

4. A message displays as shown in [Figure 51](#) requesting confirmation that you are deleting failed Vdisks and DR groups, and asks if you wish to continue. Click **OK**.




Figure 51: Message confirming Vdisk and DR group deletion

5. An informational screen similar to [Figure 52](#) is displayed and lists affected DR groups requiring deletion.




Figure 52: Command View EVA manual deletion message

6. Click an affected DR group. A message is displayed (see [Figure 53](#)) to inform you that a DR group is being deleted.

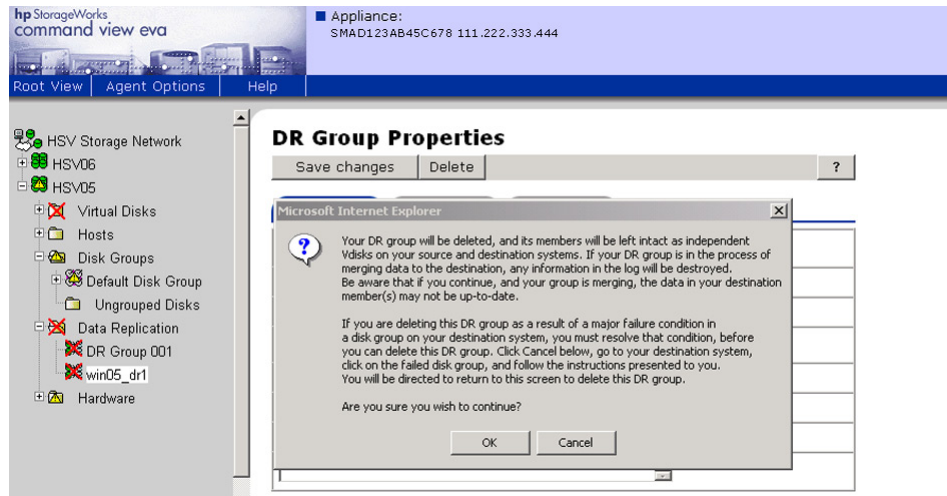


Figure 53: Message confirming DR group deletion

7. Click **OK**. The affected DR groups are deleted.
8. Click failed Vdisks that were members of the affected DR groups. A message is displayed (Figure 54) while Vdisks are being deleted.

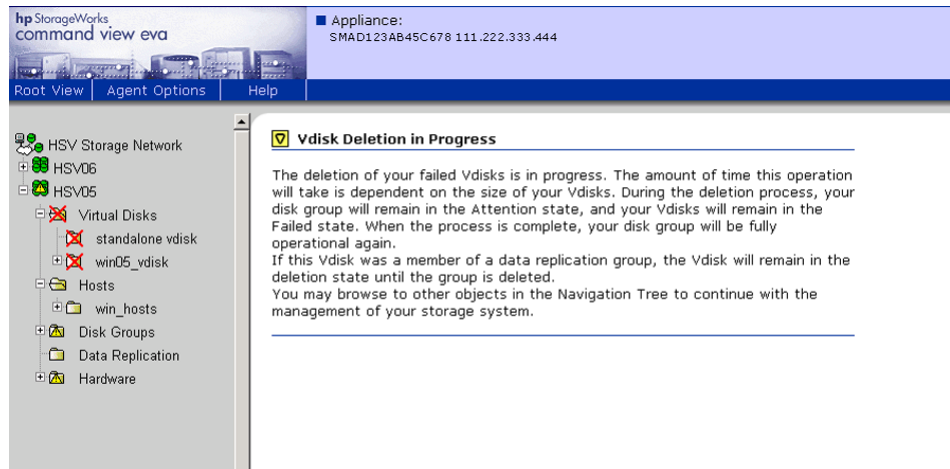


Figure 54: Vdisk Deletion in Progress message

9. When the deletion process is complete, a Command View EVA Vdisk Folder Properties screen is displayed showing the Vdisk was deleted from the tree in the left pane.
10. Repair your inoperative disk group.
11. Navigate to the surviving storage system and re-create the DR groups and set up host presentation on the repaired storage system. After normalization between the source and destination storage systems occurs, you can fail over to the repaired storage system using the procedure “Planned failover” on page 119.

Recovery when data replication was logging before failure

If data is logging when a source disk group hardware failure occurs, then the data on the destination storage system is stale (not current). You have several options:

- Use the procedure “Recovery when data replication was normal before failure” on page 143 and recover using the stale data on the destination storage system.
- Recover from a known, good point using a backup.
 - If you want to perform a failover to quickly activate the destination storage system before repairing the inoperative disk group, use the procedure “Recovery when data replication was normal before failure” on page 143, and then restore from a backup.
 - If you want to repair the inoperative disk group first, perform the repair, delete the inoperative DR groups and Vdisks on the failed system, re-create your Vdisks and DR groups, and then restore your data from an external backup.

Disk group hardware failure on the destination storage system

This section describes how to recover from an inoperative disk group on your destination storage system. Your first indications that a disk group has become inoperative may be screens like those shown in [Figure 48](#) or [Figure 49](#), except that your destination disk group status is Unknown. To recover from this condition, perform the following steps.

1. Using Command View EVA, navigate to an affected DR group on your source storage system and then click **Delete** on the DR Group Properties page ([Figure 55](#)). A message is displayed that states that your DR group will be deleted (see [Figure 56](#)) and asks if you wish to continue.

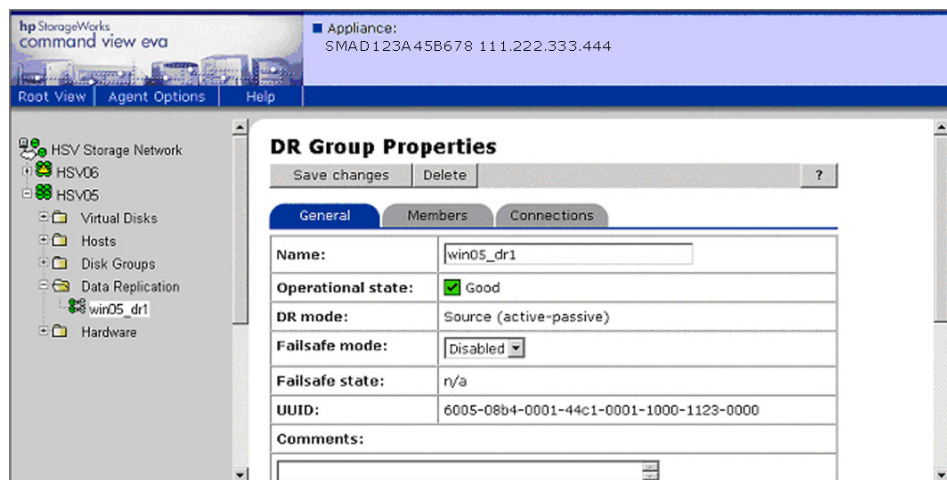


Figure 55: DR group Properties page



Figure 56: DR group deletion message

2. Click **OK**. The DR group is deleted.
3. Apply steps 1 and 2 to any other affected DR groups.
4. Navigate to the destination storage system and open the failed disk group in the Disk Group folder. You are presented with a Disk Group Hardware Failure page like that shown in [Figure 57](#).




Figure 57: Disk group Hardware Failure page

5. Click the **Start deletion process** tab. A confirmation message asks if you wish to continue by deleting the Vdisks and DR groups that have failed as a result of the hardware problem.
6. Click **OK**.
7. Navigate to each failed DR group and delete the DR groups. This deletes all Vdisks that were members of the affected DR groups.

Note: The Vdisk deletion process may take some time to conclude. Do not proceed to the next step until all the affected disks have been deleted.

8. Repair the affected disk groups.
9. Navigate to the source storage system and recreate your DR groups.
10. Navigate to the destination storage system and set up presentations for your created Vdisks.
11. With the Continuous Access user interface, click the **Refresh** icon, choose **Discover**, and click **OK**.

Troubleshooting with Multiple Sites



This chapter provides troubleshooting guidance for storage systems and links between multiple sites. The following topics are discussed:

- [Troubleshooting storage problems](#), page 153
- [Troubleshooting intersite link problems](#), page 175

Troubleshooting storage problems

With multiple sites, the decision to do a failover, as well as deciding which failover procedure to perform, can be complicated. With Continuous Access EVA V1.1, each storage system is allowed two replication relationships. This means that a storage system may have source or destination Vdisks in DR groups that replicate to as many as two other storage systems.

[Table 9](#) lists nine different replicating storage configurations with each called a “situation.” The storage system marked X is the storage system in question, or receiving the troubleshooting actions. Storage systems marked A or B are remote storage systems in a replication relationship with storage system X. A remote storage system is depicted with dotted lines. Vdisks in each storage system are marked either S (for source) or D (for destination). Match up the situation closest to the environment you want to troubleshoot and you are given a page location for troubleshooting steps located in [Table 10](#).

Table 9: Identifying your troubleshooting situation

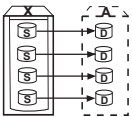
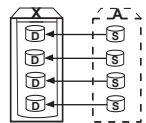
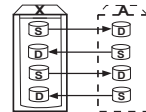
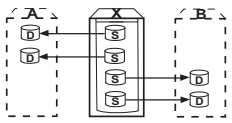
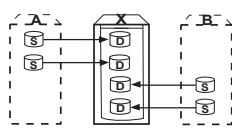
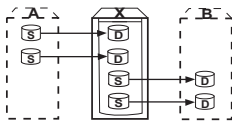
Graphical depiction of your environment	Situation description	Location
	Situation 1. One source storage system with one relationship.	page 156
	Situation 2. One destination storage system with one relationship.	page 158
	Situation 3. A storage system has DR groups with one bidirectional relationship to one other storage system.	page 160
	Situation 4. A source storage system has DR groups with replication to two destination storage systems.	page 162
	Situation 5. A destination storage system has DR groups with replication from two other storage systems.	page 164
	Situation 6. A storage system has DR groups with replication to a source storage system and from a destination storage system.	page 166

Table 9: Identifying your troubleshooting situation (Continued)

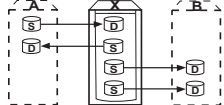
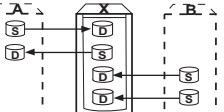
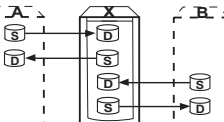
Graphical depiction of your environment	Situation description	Location
	<p>Situation 7. A storage system has DR groups with source and destination Vdisks to another storage system (site A) and only source Vdisks to a second storage system (site B).</p>	<p>page 168</p>
	<p>Situation 8. A storage system has DR groups with source and destination Vdisks to another storage system (site A) and only destination Vdisks to a second storage system (site B).</p>	<p>page 170</p>
	<p>Situation 9. A storage system has DR groups with source and destination Vdisks to two other storage systems (sites A and B).</p>	<p>page 172</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships

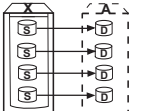
Your site situation and the troubleshooting flow	
<p>Situation 1. One source storage system with one relationship.</p>	
<p>Is the site operational (SMAs, hosts, and storage systems)?</p> <p>Yes. Are all DR groups logging and/or failsafe-locked?</p> <p>Yes: Is the destination site operational?</p> <p>Yes: Are the DR groups suspended?</p> <p>Yes: Resume all DR groups.</p> <p>No. There is a communication failure between the storage systems. You have two choices:</p> <ol style="list-style-type: none"> 1. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. If any DR groups are failsafe-locked, see "Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)" on page 132. 2. You may failover to the destination site. See "Unplanned failover" on page 127. <p>No. Repair the destination site or replace with new hardware. Use "Return operations to replaced new storage hardware" on page 134 or "Disk group hardware failure on the destination storage system" on page 150.</p> <p>No. Some DR groups are still replicating. If you cannot fix the problem with DR groups, you may want to do a failover from the other storage system until the problem is resolved. See "Planned failover" on page 119 for operational DR groups or "Unplanned failover" on page 127 for failed DR groups.</p> <p>No. Are there problems with the current SMA managing storage system X?</p> <p>Yes. Use another SMA. See "Moving storage management to another SMA" on page 93.</p>	

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow
<p>No. Continue to use this SMA. Are you having hardware failures on your storage?</p> <p>Yes. Is your storage still operating through redundant components?</p> <p>Yes. Repair your hardware. Refer to EVA documentation.</p> <p>No. See "Unplanned failover" on page 127.</p> <p>No. Are the hosts operational and able to access the storage?</p> <p>Yes. If the recovery procedures do not address your issue, refer to other documentation or contact technical support.</p> <p>No. See "Unplanned failover" on page 127.</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

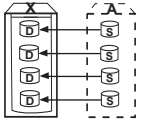
Your site situation and the troubleshooting flow	
<p>Situation 2. One destination storage system with one relationship.</p>	
<p>Is the site operational (SMAs, hosts, and storage systems)?</p> <p>Yes. Have all DR groups lost connection with the source?</p> <p>Yes: Is the source site operational?</p> <p>Yes: Are the DR groups suspended?</p> <p>Yes: Resume all DR groups.</p> <p>No. There is a communication failure between the storage systems. You have two choices:</p> <ol style="list-style-type: none"> 1. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132. 2. You may failover to this site. See “Unplanned failover” on page 127. <p>No. Proceed with “Unplanned failover” on page 127 at this storage system.</p> <p>No. Some DR groups are still replicating. If you cannot fix the problem with DR groups, you may want to do a failover from this storage system until the problem is resolved. See “Planned failover” on page 119 for operational DR groups or “Unplanned failover” on page 127 for failed DR groups. Repair your hardware (refer to EVA documentation).</p> <p>No. Are there problems with the current SMA managing storage system X?</p> <p>Yes. Use another SMA. See “Moving storage management to another SMA” on page 93.</p> <p>No. Continue to use this SMA. Are you having hardware failures on your storage?</p>	

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow
<p>Yes. Is your storage still operating through redundant components?</p> <p>Yes. Repair your hardware. Refer to EVA documentation.</p> <p>No. Repair your hardware. Refer to EVA documentation. Determine if the source site is still operational and logging (situation 1) and that no further action is needed.</p> <p>No. If the recovery procedures do not address your issue, refer to other documentation or contact technical support.</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow	
Situation 3. A storage system has DR groups with one bidirectional relationship to one other storage system.	
<p>Is the site operational (SMAs, hosts, and storage systems)?</p> <p>Yes. Are all DR groups logging on storage system X?</p> <p>Yes: Is the other site operational?</p> <p>Yes: Are the DR groups suspended?</p> <p>Yes: Resume all DR groups.</p> <p>No. There is a communication failure between the storage systems. You have three choices:</p> <ol style="list-style-type: none"> 1. You may continue logging at this site. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. 2. You may continue logging at this site and failover all DR groups on destination storage system A to this site. See "Unplanned failover" on page 127. 3. You may continue logging at the other site and failover all DR groups on destination storage system X to the other site. See "Unplanned failover" on page 127. <p>No. Failover all DR groups on the destination storage system to this site. See "Unplanned failover" on page 127.</p> <p>No. The DR groups are still replicating. If you cannot fix the problem with DR groups, you may want to do a failover from the other storage system until the problem is resolved. See "Planned failover" on page 119 for operational DR groups or "Unplanned failover" on page 127 for failed DR groups. Repair your hardware (refer to EVA documentation).</p> <p>No. Are there problems with the current SMA managing storage system X?</p> <p>Yes. Use another SMA. See "Moving storage management to another SMA" on page 93.</p> <p>No. Continue to use this SMA. Are you having hardware failures on your storage?</p>	

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow
<p>Yes. Is your storage still operating through redundant components?</p> <p>Yes. Repair your hardware. Refer to EVA documentation.</p> <p>No. Perform a failover. See “Unplanned failover” on page 127.</p> <p>No. If the recovery procedures do not address your issue, refer to other documentation or contact technical support.</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow	
Situation 4. A source storage system has DR groups with replication to two destination storage systems.	<p>The diagram illustrates a source storage system (S) with two DR groups, A and B. Group A replicates to destination D1, and Group B replicates to destination D2. Arrows indicate the replication flow from the source to the destinations.</p>
<p>Is the site operational (SMAs, hosts, and storage systems)?</p>	<p>Yes. Are all DR groups to one destination storage system logging on the source storage system?</p> <p>Yes: Is each destination site operational?</p> <p>Yes: Are the DR groups suspended?</p> <p>Yes: Resume all DR groups.</p> <p>No. There is a communication failure between the storage systems. You have two choices:</p> <ol style="list-style-type: none"> 1. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132. 2. You may failover to a destination site. See “Unplanned failover” on page 127. <p>No. Repair the destination site or replace with new hardware. Use “Return operations to replaced new storage hardware” on page 134 or “Disk group hardware failure on the destination storage system” on page 150.</p> <p>No. The DR groups are still replicating. If you cannot fix the problem with DR groups, you may want to do a failover from the other storage system until the problem is resolved. See “Planned failover” on page 119 for operational DR groups or “Unplanned failover” on page 127 for failed DR groups.</p> <p>No. Are there problems with the current SMA managing storage system X?</p> <p>Yes. Use another SMA. See “Moving storage management to another SMA” on page 93.</p> <p>No. Continue to use this SMA. Are you having hardware failures on your storage?</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow
<p>Yes. Is your storage still operating through redundant components?</p> <p>Yes. Repair your hardware. Refer to EVA documentation.</p> <p>No. Perform a failover. See "Unplanned failover" on page 127.</p> <p>No. If the recovery procedures do not address your issue, refer to other documentation or contact technical support.</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

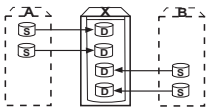
Your site situation and the troubleshooting flow	
<p>Situation 5. A destination storage system has DR groups with replication from two other storage systems.</p>	
<p>Is the site operational (SMAs, hosts, and storage systems)?</p> <p>Yes. Have all DR groups on this destination storage system lost connection with one of the source storage systems?</p> <p>Yes: Is each source site operational?</p> <p>Yes: Are the DR groups suspended?</p> <p>Yes: Resume all DR groups.</p> <p>No. There is a communication failure between two of the storage systems. You have two choices:</p> <ol style="list-style-type: none"> 1. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132. 2. You may failover to this site. See “Unplanned failover” on page 127. <p>No. Proceed with “Unplanned failover” on page 127.</p> <p>No. Some DR groups are still replicating. If you cannot fix the problem with DR groups, you may want to do a failover from the other storage system until the problem is resolved. See “Planned failover” on page 119 for operational DR groups or “Unplanned failover” on page 127 for failed DR groups.</p> <p>No. Are there problems with the current SMA managing storage system X?</p> <p>Yes. Use another SMA. See “Moving storage management to another SMA” on page 93.</p> <p>No. Continue to use this SMA. Are you having hardware failures on your storage?</p>	

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow
<p>Yes. Is your storage still operating through redundant components?</p> <p>Yes. Repair your hardware. Refer to EVA documentation.</p> <p>No. Repair your hardware. Refer to EVA documentation. Determine if the source site is still operational and logging (situation 1) and that no further action is needed.</p> <p>No. If the recovery procedures do not address your issue, refer to other documentation or contact technical support.</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow	
<p>Situation 6. A storage system has DR groups with replication to a source storage system and from a destination storage system.</p>	
<p>Is the site operational (SMAs, hosts, and storage systems)?</p> <p>Yes. Answer these two questions:</p> <ol style="list-style-type: none"> Are all source DR groups on storage system X logging and/or failsafe-locked. <p>Yes. Is the site B operational?</p> <p>Yes: Are the DR groups suspended?</p> <p>Yes: Resume all DR groups.</p> <p>No. There is a communication failure between storage systems X and B. You have two choices:</p> <ol style="list-style-type: none"> If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132. You may failover to this site. See “Unplanned failover” on page 127. <p>No. Repair the destination site or replace with new hardware. Use “Return operations to replaced new storage hardware” on page 134 or “Disk group hardware failure on the destination storage system” on page 150.</p> <p>No. Some DR groups are still replicating between sites X and B. If you cannot fix the problem with DR groups, you may want to do a failover from the other storage system until the problem is resolved. See “Planned failover” on page 119 for operational DR groups or “Unplanned failover” on page 127 for failed DR groups.</p>	

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow
<p>2. Have all destination DR groups on this storage system lost their connections with the source DR groups on site A?</p> <p>Yes. Is site A operational?</p> <p>Yes. There is a communication failure between the storage systems A and X. You have two choices:</p> <p>A. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows.</p> <p>If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132.</p> <p>B. You may failover to this site. See “Unplanned failover” on page 127.</p> <p>No. Proceed with “Unplanned failover” on page 127.</p> <p>No. Some DR groups are still replicating between sites X and A. If you cannot fix the problem with DR groups, you may want to do a failover from the other storage system until the problem is resolved. See “Planned failover” on page 119 for operational DR groups or “Unplanned failover” on page 127 for failed DR groups.</p> <p>No. Are there problems with the current SMA managing storage system X?</p> <p>Yes. Use another SMA. See “Moving storage management to another SMA” on page 93.</p> <p>No. Continue to use this SMA. Are you having hardware failures on your storage?</p> <p>Yes. Is your storage still operating through redundant components?</p> <p>Yes. Repair your hardware. Refer to EVA documentation.</p> <p>No. Failover to site B and continue logging at site A. See “Unplanned failover” on page 127.</p> <p>No. If the recovery procedures do not address your issue, refer to other documentation or contact technical support.</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow	
<p>Situation 7. A storage system has DR groups with source and destination Vdisks to another storage system (site A) and only source Vdisks to a second storage system (site B).</p>	<p>The diagram illustrates a storage system X (represented by a vertical rectangle) connected to two other storage systems, A and B, which are enclosed in dashed boxes. Storage system X has four Vdisk icons (labeled S and D) arranged vertically. Site A has two Vdisk icons (S and D) with arrows pointing from X to A. Site B has two Vdisk icons (S and D) with arrows pointing from X to B. This represents a configuration where site X has DR groups with source and destination Vdisks to site A, and only source Vdisks to site B.</p>
<p>Is the site operational (SMAs, hosts, and storage systems)?</p>	<p>Yes. Are all DR groups on storage system X, with source Vdisks to either storage systems A or B, logging?</p> <p>Yes: Are the DR groups suspended?</p> <p>Yes: Resume all DR groups.</p> <p>No. There is a communication failure between the storage systems. Answer these two questions:</p> <ol style="list-style-type: none"> Is site A operational? <ul style="list-style-type: none"> Yes. There is a communication failure between storage systems A and X. There are three choices: <ol style="list-style-type: none"> You may continue logging at all sites. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132. Stop using site X. You may continue logging at sites A and B, and failover all applicable DR groups to sites A and B. See “Unplanned failover” on page 127. If any DR groups are failsafe-locked at sites A, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132. Stop using site A. Continue logging at this site. If the DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132. No. Failover all DR groups where the source Vdisks are now at site A to this site. See “Unplanned failover” on page 127.

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow
<p>2. Is site B operational?</p> <p>Yes. There is a communication failure between the storage systems X and B. You have two choices:</p> <p>A. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows.</p> <p>If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132.</p> <p>B. You may failover to site X. See “Planned failover” on page 119.</p> <p>No. Repair the destination site or replace with new hardware. Use “Return operations to replaced new storage hardware” on page 134 or “Disk group hardware failure on the destination storage system” on page 150.</p> <p>No. Some DR groups are still replicating. If you cannot fix the problem with DR groups, you may want to do a failover from the other storage system until the problem is resolved. See “Planned failover” on page 119 for operational DR groups or “Unplanned failover” on page 127 for failed DR groups.</p> <p>No. Are there problems with the current SMA managing storage system X?</p> <p>Yes. Use another SMA. See “Moving storage management to another SMA” on page 93.</p> <p>No. Continue to use this SMA. Are you having hardware failures on your storage?</p> <p>Yes. Is your storage still operating through redundant components?</p> <p>Yes. Repair your hardware. Refer to EVA documentation.</p> <p>No. Failover to site B and continue logging at site A. See “Unplanned failover” on page 127.</p> <p>No. If the recovery procedures do not address your issue, refer to other documentation or contact technical support.</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

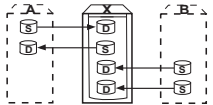
Your site situation and the troubleshooting flow	
<p>Situation 8. A storage system has DR groups with source and destination Vdisks to another storage system (site A) and only destination Vdisks to a second storage system (site B).</p>	
<p>Is the site operational (SMAs, hosts, and storage systems)?</p> <p>Yes. Are all DR groups with source Vdisks, that replicate to another storage system, logging?</p> <p>Yes: Are the DR groups suspended?</p> <p>Yes: Resume all DR groups.</p> <p>No. There is a communication failure between storage systems A and X. Answer these two questions:</p> <ol style="list-style-type: none"> Is site A operational? <p>Yes. There is a communication failure between the storage systems. There are three choices:</p> <ol style="list-style-type: none"> You may continue logging at all sites. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. No action required. If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132. Stop using site X. Continue logging at sites A and B. Failover all applicable DR groups to sites A and B. See “Unplanned failover” on page 127. <p>If any DR groups are failsafe-locked at sites A or B, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132.</p> <ol style="list-style-type: none"> Stop using site A. Perform a failover on all applicable DR groups. See “Unplanned failover” on page 127. No. Failover all DR groups with destination Vdisks at site A to this site. See “Unplanned failover” on page 127. 	

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow
<p>2. Is site B operational?</p> <p>Yes. There is a communication failure between storage systems B and X. You have two choices:</p> <p>A. You may continue logging at site B. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. No action required.</p> <p>If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132.</p> <p>B. Stop using site B. See “Unplanned failover” on page 127.</p> <p>If any DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. No action required.</p> <p>If the DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132.</p> <p>No. Perform “Unplanned failover” on page 127.</p> <p>No. Some DR groups are still replicating. If you cannot fix the problem with DR groups, you may want to do a failover from the other storage system until the problem is resolved. See “Planned failover” on page 119 for operational DR groups or “Unplanned failover” on page 127 for failed DR groups.</p> <p>No. Are there problems with the current SMA managing storage system X?</p> <p>Yes. See “Moving storage management to another SMA” on page 93.</p> <p>No. Continue to use this SMA. Are you having hardware failures on your storage?</p> <p>Yes. Is your storage still operating through redundant components?</p> <p>Yes. Repair your hardware. Refer to EVA documentation.</p> <p>No. Perform “Unplanned failover” on page 127 to sites A and B.</p> <p>No. If the recovery procedures do not address your issue, refer to other documentation or contact technical support.</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

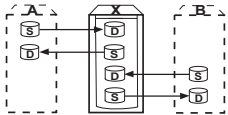
Your site situation and the troubleshooting flow	
<p>Situation 9. A storage system has DR groups with source and destination Vdisks to two other storage systems (sites A and B).</p>	
<p>Is the site operational (SMAs, hosts, and storage systems)?</p> <p>Yes. Are all source Vdisks logging or failsafe-locked?</p> <p>Yes: Are the DR groups suspended?</p> <p>Yes: Resume all DR groups.</p> <p>No. There is a communication failure between the storage systems.</p> <p>Answer these two questions:</p> <ol style="list-style-type: none"> Is site A operational? <p>Yes. There is a communication failure between storage systems X and A. There are three choices:</p> <ol style="list-style-type: none"> You may continue logging at all sites. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. No action required. If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132. Stop using site X. You may continue logging at sites A and B, and failover all applicable DR groups to sites A and B. See “Unplanned failover” on page 127. <p>If any DR groups are failsafe-locked at sites A or B, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132.</p> <ol style="list-style-type: none"> Stop using site A. Perform a failover on all applicable DR groups. See “Unplanned failover” on page 127. No. Failover all DR groups with destination Vdisks at site A to this site. See “Unplanned failover” on page 127. 	

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow
<p>2. Is site B operational?</p> <p>Yes. There is a communication failure between storage systems X and B. You have three choices:</p> <p>A. You may continue logging at all sites. If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. No action required.</p> <p>If any DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132.</p> <p>B. Stop using site X. You may continue logging at sites A and B, and failover all DR groups to sites A and B. See “Unplanned failover” on page 127.</p> <p>If the DR groups are logging, continue logging. When the connection comes up, the log on the source DR group will merge with the destination or perform a full copy if the log overflows. No action required.</p> <p>If the DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132.</p> <p>C. Stop using site B by performing a failover on all applicable DR groups. See “Unplanned failover” on page 127. Continue logging at your site. If the DR groups are failsafe-locked, see “Resumption of operations if unable to access destination while source in failsafe-locked state (extended period of time)” on page 132.</p> <p>No. Failover all DR groups at the destination storage system (site B) to your site. See “Unplanned failover” on page 127.</p> <p>No. Some DR groups are still replicating. If you cannot fix the problem with DR groups, you may want to do a failover from the other storage system until the problem is resolved. See “Planned failover” on page 119 for operational DR groups or “Unplanned failover” on page 127 for failed DR groups.</p>

Table 10: Troubleshooting a site storage system with single or multiple site replicating relationships (Continued)

Your site situation and the troubleshooting flow
<p>No. Are there problems with the current SMA managing storage system X?</p> <p>Yes. See "Moving storage management to another SMA" on page 93.</p> <p>No. Continue to use this SMA. Are you having hardware failures on your storage?</p> <p>Yes. Is your storage still operating through redundant components?</p> <p>Yes. Repair your hardware. Refer to EVA documentation.</p> <p>No. Perform "Unplanned failover" on page 127 to sites A and B.</p> <p>No. If the recovery procedures do not address your issue, refer to other documentation or contact technical support.</p>

Troubleshooting intersite link problems

If you suspect a problem with your intersite link, HP recommends that you first log into the switches at each end of the link. You can access a switch from either a Web display or a Telnet session. Once the switch is accessed, you can obtain the current status of the switch ports. For example, with B-series switches the command is `switchshow`. Look for the identity of all interswitch and intersite links. If the links are active, obtain the current performance information (the command with B-series switches is `portperfshow`). Between the two displays you should be able to determine if the fabric is intact, and if so, that data is moving through the link. If not, then the link is down.

The rate at which data is moving across the links should be consistent with the rates of data written to a storage system. For example, if three applications are writing data totaling 10 MB/sec, and all the data must be replicated, then there should be 10 MB/sec of data moving through the intersite links.

Continuous Access EVA Support Procedures

8

This chapter describes best practice procedures that are used in support of Continuous Access EVA.

This chapter covers the following topics:

- [Creating a destination Snapclone before making a full copy](#), page 177
- [Data movement using a Snapclone](#), page 178
- [Three-site cascaded data replication using Snapclones](#), page 180
- [Bootless DR group planned failover with Linux using LVM in standalone mode or with SuSE SLES 8 running LifeKeeper 4.4.3](#), page 197

Creating a destination Snapclone before making a full copy

Note: A Business Copy license is required for the following procedure.

When logging occurs on a source storage system, a temporary disparity is created between data being processed at the local site and the data that exists at the remote location. A merge or full copy corrects this disparity later when the reason for the interruption is remedied. A merge sends data from the log disk in write order so it remains crash consistent. However, this is not the case with a full copy.

When a log fills to the point where it is marked for a full copy, there is a risk to the destination copy of the data once the process begins. This risk is due to the nature of a full copy, which copies data in 1-MB chunks starting at the beginning of the source Vdisk. This full-copy process does not maintain write ordering, and if it is not completed due to a second error, such as a loss of the source storage system, it leaves the destination storage system in an indeterminate state. Therefore, to prevent loss of data, best practice suggests creating a Snapclone of destination Vdisks containing critical or important data prior to starting a full copy. If a major failure then occurs at the local site during a full copy, the Snapclone provides a

clean copy of data as it existed before full copy writes were started to the destination storage system. However, any new writes that occurred on the source between the time the Snapclone was created and the major failure would result in the loss of the new writes.

The following procedure describes the steps to take in a situation where you lose the connection between a source and destination storage system and want to protect against a second failure when performing a full copy. Best practice suggests the creation of a destination Snapclone whenever the link outage is expected to last more than several minutes. However, you cannot use this procedure if a full copy has been started.

1. Using Command View EVA or the Continuous Access user interface, navigate to each affected DR group and suspend replication.
2. When able, use the managing SMA to make a Snapclone of the destination Vdisks, using the procedures described in HP StorageWorks Business Copy documentation.
3. Using Command View EVA or the Continuous Access user interface, navigate to each affected DR group and resume replication. This will only enable replication if the links are still down.

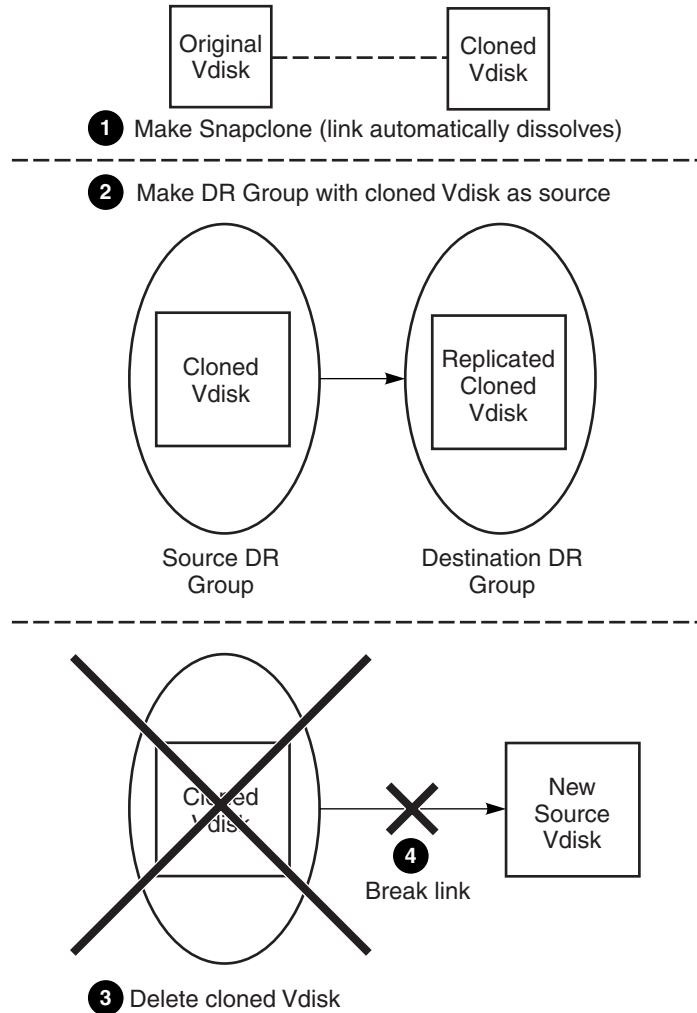
Data movement using a Snapclone

Note: A Business Copy license is required for the following procedure.

This procedure can be used to move a copy of your data residing on a Vdisk to a remote location by the use of a Snapclone. A Snapclone is an exact copy of your Vdisk at the particular point-in-time it was created. The Vdisk being copied to the remote site becomes part of a DR group that can then be used as a new source Vdisk. This procedure can be used with several data movement services, such as:

- **Data distribution**—Pushing copies of data to other geographic locations to make it locally accessible.
- **Data migration**—Moving data to a new location or to one with a larger storage capacity.

Figure 58 provides a high-level summary of the following steps that perform data movement using a Snapclone:



CXO8068A

Figure 58: Creating a DR group from a Snapclone

1. Make a Snapclone of the Vdisk containing the data to be moved. Refer to the HP StorageWorks Business Copy documentation for procedures on creating Snapclones.

After the Snapclone is created, the link from the Snapclone to its original Vdisk dissolves, and the Snapclone becomes a separate Vdisk.

2. Create a DR group with the new Snapclone-created Vdisk linked to the remote storage system where you want the data to reside. The creation of the DR group will replicate the Vdisk to your desired destination. For procedures on creating a DR group, refer to “Creating DR Groups” in Chapter 4.
3. Delete the source-cloned Vdisk. You have the option of keeping the remote Vdisk or deleting the remote Vdisk
4. Choose to keep remote Vdisk.

The data now resides as a new Vdisk on the remote storage system. It can be used as a source for another DR group, subject to the restriction that a storage system can be involved in a replicating relationship with only one other storage system.

Three-site cascaded data replication using Snapclones

Note: A Business Copy license is required for the following procedure.

This procedure describes how to move copies of your data to a second remote location with Command View EVA and the use of Snapclones. The remote location can be a storage system without a replicating relationship to the array where the data was created. Exact copies of the Vdisks containing the data are created by using Snapclones, and these are placed into a DR group for movement to the remote system.

For example, a production environment contains a DR group that replicates between storage systems HSV05 and HSV06 (see [Figure 59](#)). The DR group contains two Vdisks (05-06vdisk1 and 05-06vdisk2) that are to be archived on another storage system (HSV18). A Snapclone of each Vdisk will be created on the destination storage system (HSV06). After presentation to a phantom host (a non-existent host set up just for presentation purposes only, but is required for the creation of a DR group), these members will be added to a DR group called DR Snapclone1. This DR group now resides on a source storage system that replicates

to the desired destination storage system (HSV18). At the remote location, the Vdisk members can be removed from the DR group, renamed, and archived. The following procedure describes one way this can be done:

Note: For this procedure, HSV05 is called the source storage system, HSV06 (the destination for the DR group from HSV05) is called the intermediate storage system, and HSV18 is referred to as the remote storage system.

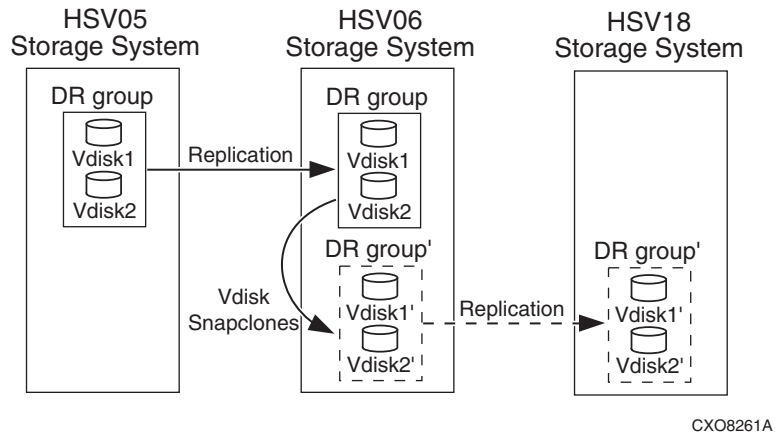


Figure 59: Data movement using Snapclones example

1. HP recommends that any DR group containing more than one copy set have failsafe mode enabled. To enable failsafe mode:
 - a. In the navigation pane, expand the desired storage system and select the **Data Replication** folder. Any DR groups in this folder are displayed.
 - b. Select the desired DR group. A DR Group Properties pane is displayed (Figure 60).
 - c. Choose the **General** tab.
 - d. In the **Failsafe mode** drop-down list, select **Enabled**.
 - e. Choose the **Save changes** tab.




Figure 60: Setting a DR group to failsafe mode

2. HP recommends that DR groups in this procedure be set for synchronous write mode. To enable synchronous write mode:
 - a. In the navigation pane, expand the desired storage system and select the **Data Replication** folder. Any DR groups in this folder are displayed.
 - b. Select the desired DR group. A DR Group Properties pane is displayed (Figure 61).
 - c. Choose the **Connections** tab.
 - d. In the **Write mode** drop-down list, select **Synchronous**.
 - e. Choose the **Save changes** tab.
3. If normalization is occurring to members of the DR group to be moved, wait for the members to normalize. Normalization is complete when the copy state of the members listed in the Group Members pane on the DR Group Properties page show to be normal (see Figure 62). If an application requires that I/O be suspended before creation of a Snapclone, suspend I/O at this time.




Figure 61: Setting a DR group for synchronous replication




Figure 62: Checking normalization of DR group members

4. Create a Snapclone of each Vdisk on the intermediate storage system (HSV06).
 - a. A host presentation is required when creating a Snapclone, but it does not have to be a physical host containing FCAs. A phantom host can be created as described below to act as a presented host.
 - 1) In the Navigation pane, select the Hosts folder where you want to create a host. The Host Folder Properties page is displayed (Figure 63).
 - 2) Click the **Add host** tab. The Add a Host page is displayed (Figure 64).




Figure 63: Host Folder Properties page

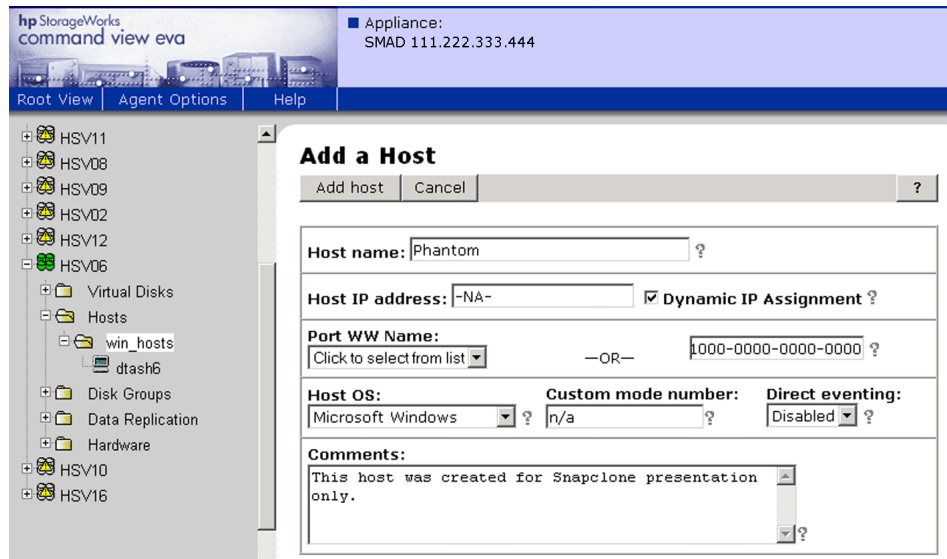


Figure 64: Add a Host page

- 3) Enter a host name in the **Host name** field and a fictitious WWN in the box to the right of the Port WW Name field. In this example, a port WW Name of 1000-0000-0000-0000 is used. Click the **Add host** tab. An Operation completed page is displayed ([Figure 65](#)).




Figure 65: Operation succeeded page.

- 4) Click **OK**.
 - b. Select the active member of the Vdisk in the navigation pane. The Vdisk Active Member Properties page is displayed (see [Figure 66](#)).



Figure 66: Vdisk Active Member Properties page

- c. Click the **Create Snapclone** tab. The Create a Snapclone page 1 is displayed (see [Figure 67](#)).

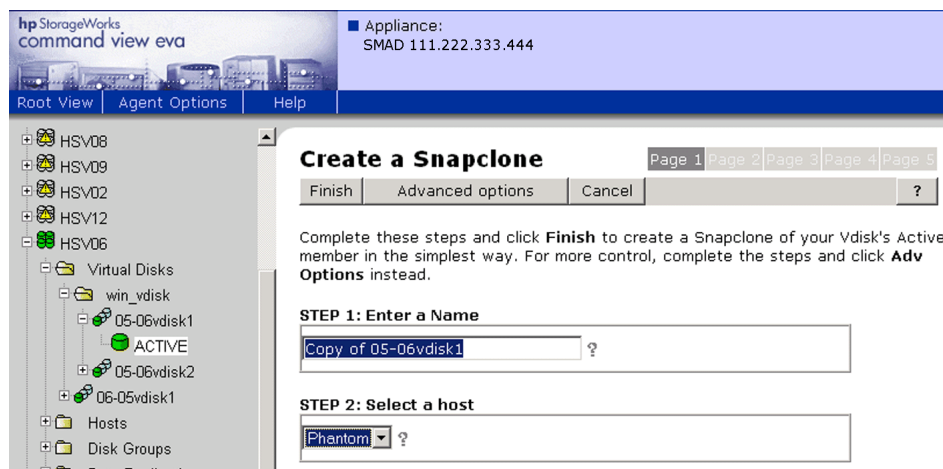


Figure 67: Create a Snapclone page 1

- d. In the **Step 1** box, enter a name for the Snapcloned Vdisk or use the default.
 - e. In the **Step 2** box, select a host for presentation. As mentioned in step a, this does not have to be a real host.
 - f. Some operating systems may require you to choose the **Advanced options** tab for additional information. In our example it is not necessary. Click the **Finish** button. A message appears asking if you want to create a Snapclone in the same DR group as the original Vdisk.
 - g. Click **OK**. The Snapclone is created. The new Snapclone will appear in the navigation pane and a page appears showing the operation succeeded. Click **OK**.
 - h. Perform steps b through g for all Vdisks in your DR group that you want to move.
 - i. If you previously suspended the application I/O, resume the I/O at this time.
 - j. If you previously set the DR group to failsafe mode, disable failsafe mode if you choose to return to this mode.
 - k. If you previously set the DR group for synchronous write mode, change back to asynchronous write mode if you choose to return to this mode.
5. Place the Snapcloned Vdisks into a new DR group.
- a. Select the active member of one of the Snapcloned Vdisks in the navigation pane. The Vdisk Active Member Properties page is displayed ([Figure 66](#)). Click the **Data Replication** tab. The Vdisk Active Member Properties page opens again with an **Add member** tab ([Figure 68](#)).

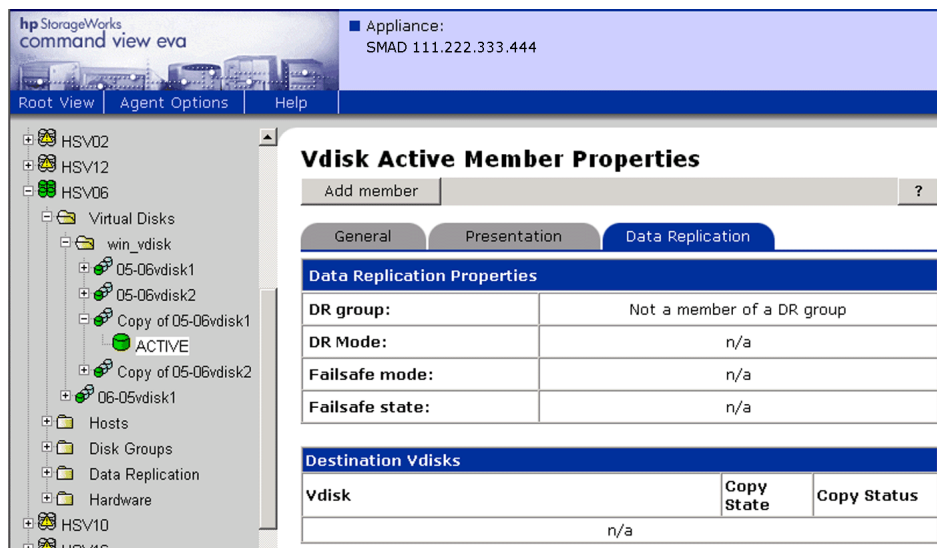


Figure 68: Vdisk Active Member Properties page with Add member tab

- b. Click the **Add member** tab. The Create a DR Group page is displayed (Figure 69).
- c. Fill in the appropriate fields to name the DR group and designate the destination storage system. In this example the DR group replicates to the remote storage system (HSV18).
- d. Click the **Create** tab. A message appears stating that a DR group will be created and that you must set up presentations as soon as possible. Click **OK**. You will receive confirmation that the DR group was successfully created. Click **OK**.
- e. Add any additional Vdisks to the DR group as described in steps 3b through 3e.
- f. Wait for normalization to complete. Normalization is complete when the copy state of the members listed in the DR Group Properties pane are shown as normal (see Figure 62).

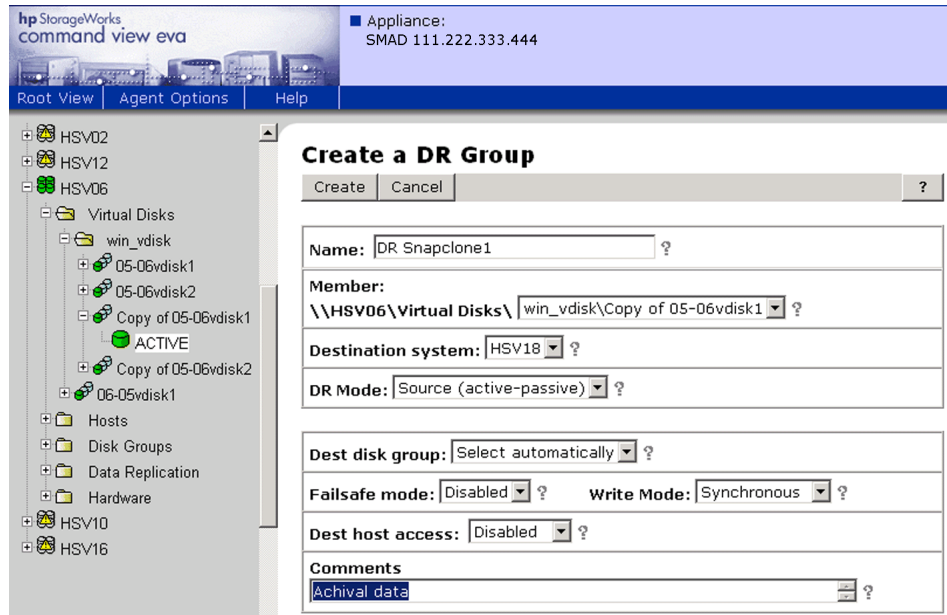


Figure 69: Create a DR Group page

6. Unpresent the host from the Snapcloned Vdisks in the DR group on the intermediate storage system (HSV06).
 - a. Click the active member of a Snapcloned Vdisk on the navigation pane. The Vdisk Active Member Properties page is displayed (Figure 66).
 - b. Click the **Presentation** tab. The Vdisk Active Member Properties page is displayed with an Unpresent tab (Figure 70).
 - c. Click the **Unpresent** tab. The Unpresent Host(s) page opens (Figure 71).

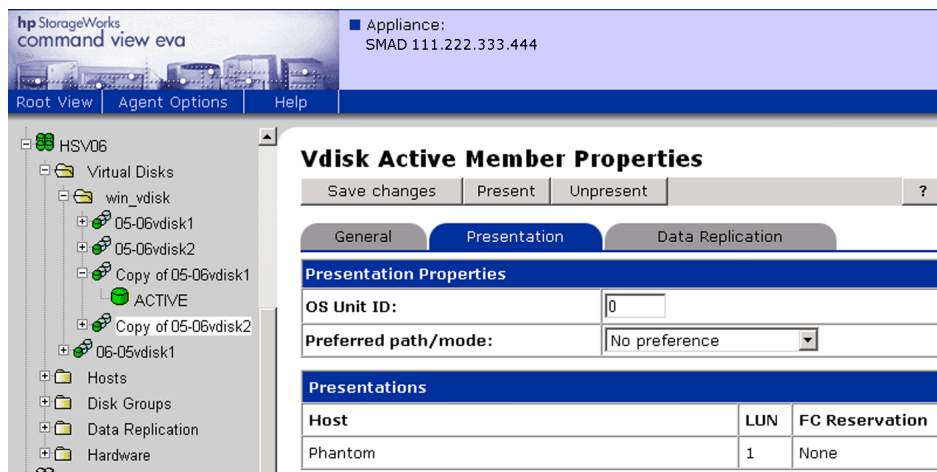


Figure 70: Vdisk Active Member Properties page with Unpresent tab

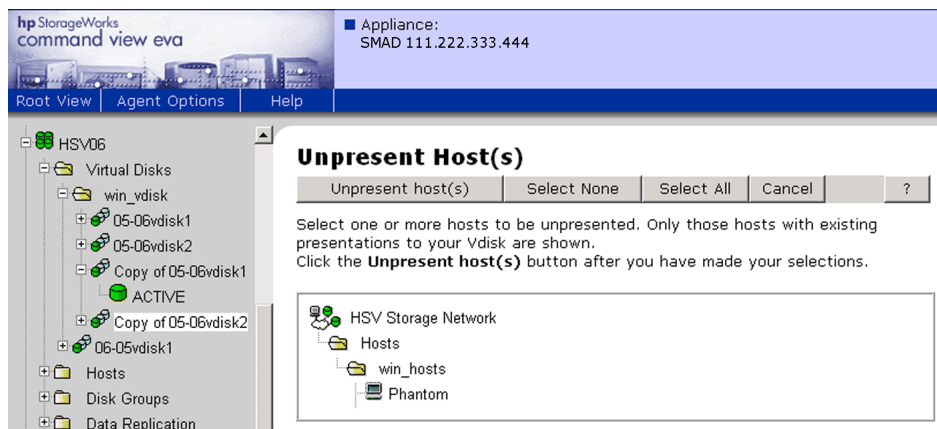


Figure 71: Unpresent Host(s) page

- d. Select the name of the host presented to the Vdisk, then choose the **Unpresent host(s)** tab. A message confirms your choice and asks if you wish to continue. Click **OK**. You will receive an Operation succeeded page to confirm the host was unpresented. Click **OK**.
- e. Unpresent the hosts to all remaining Vdisk in the DR group.

7. Remove the Snapcloned Vdisks from the intermediate storage system (HSV06).
 - a. Select the active member of a Snapcloned Vdisk in the navigation pane. The Vdisk Active Member Properties page is displayed (Figure 66).
 - b. Click the **Data Replication** tab. The Vdisk Active Member Properties page is displayed with a **Remove member** tab (Figure 72).

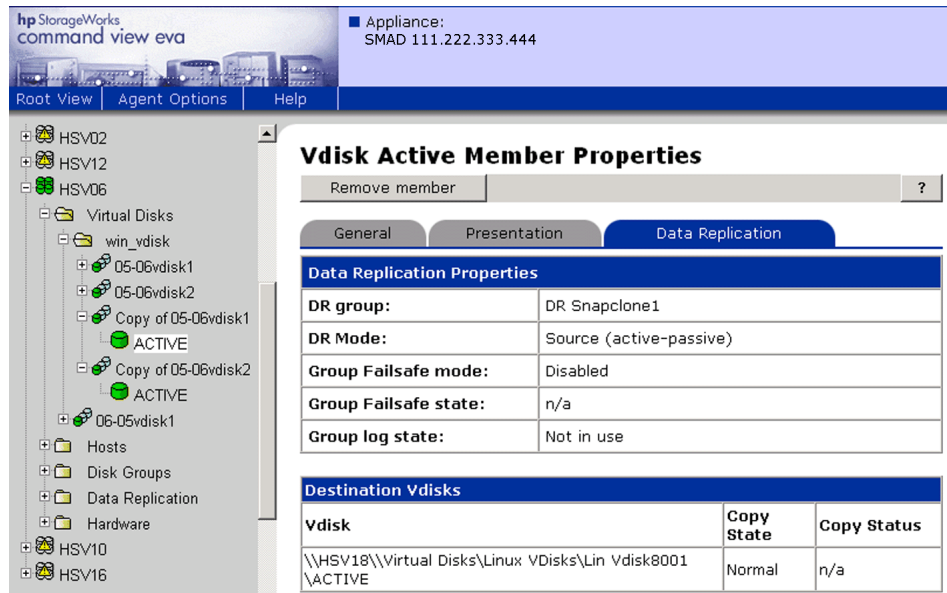


Figure 72: Vdisk Active Member Properties with Remove member tab

- c. Click the **Remove member** tab. A message confirms your choice and asks if you wish to continue. Click **OK**. A Keep/Delete Remote Mirror Vdisk page opens (Figure 73).

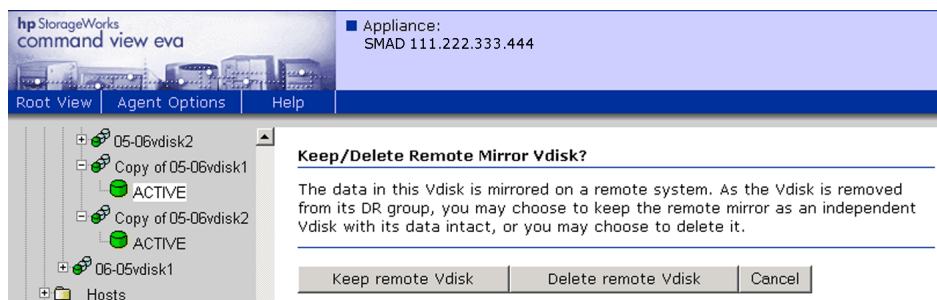


Figure 73: Keep/Delete Remote Mirror Vdisk page

- d. Choose the **Keep remote Vdisk** tab. A message appears to confirm your choice and asks if you wish to continue. Click **OK**. An Operation succeeded message is displayed. Click **OK**.
- e. Remove the remaining Snapcloned Vdisks from the DR group, leaving the remote Vdisks intact.
- f. Select a Snapcloned Vdisk name in the navigation pane. The Vdisk Family Properties window opens (Figure 74).

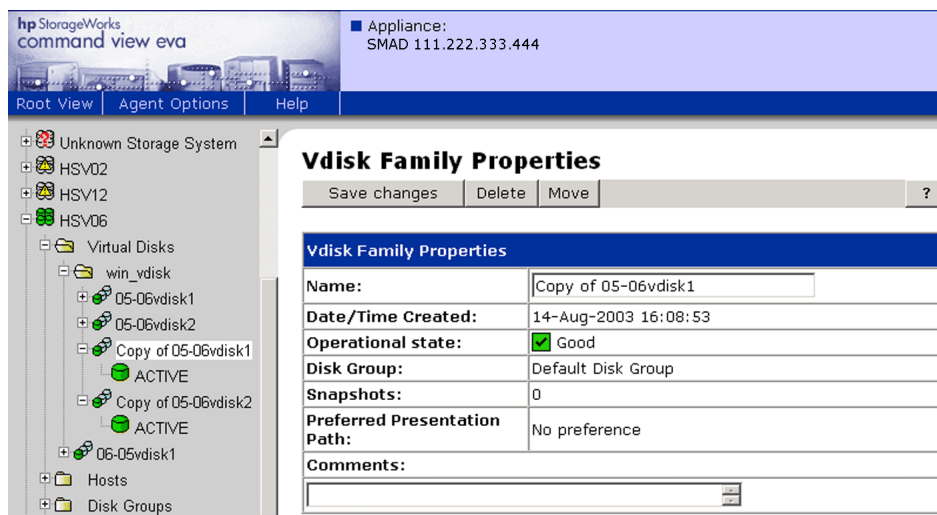


Figure 74: Vdisk Family Properties window

- g. Choose the **Delete** tab. A message appears to caution you that the Vdisk data will be lost and asks if you wish to continue. Click **OK**. An Operation succeeded page is displayed. Click **OK**.
- h. Delete the remaining Snapcloned Vdisks from the intermediate storage system (HSV06).
8. On the remote storage system (HSV18), change the write protection of the mirrored Snapcloned Vdisks to that of no write protection.
 - a. Select the active member of a mirrored Snapcloned Vdisk in the navigation pane on the remote storage system (HSV18). The Active Vdisk Member Properties page is displayed (Figure 75).

The screenshot shows the HP StorageWorks command view EVA interface. The left navigation pane displays a tree structure for HSV18, including Virtual Disks, Linux VDisks, NetWare VDisks, win_vdisk, Archive1, Archive2, Copy of 05-06vdisk1 (ACTIVE), Copy of 05-06vdisk2, Hosts, Disk Groups, Data Replication, and Hardware. The main pane displays the 'Vdisk Active Member Properties' page for the selected 'ACTIVE' member. The page has tabs for General, Presentation, and Data Replication. The General tab is active, showing fields for Name (ACTIVE), Family Name (Copy of 05-06vdisk1), World Wide LUN Name (6005-08b4-0001-0a70-0002-0000), and UUID (6005-08b4-0001-0a70-0002-0000-0048-0000). The Condition/State section shows Operational State as Good. The Date/Time section shows Created as 14-Aug-2003 16:31:34. The Cache Policies section shows Write as Mirrored write-back and Read as On. The Attributes section shows Type as Original, Disk Group as Disk Group 2, Capacity Req as 1 GB, Capacity Used as 1 GB, Redundancy as Vraid1, and Write Protect as No.

Vdisk Active Member Properties	
Save changes Create snapshot Create Snapclone ?	
General Presentation Data Replication	
Identification	
Name:	ACTIVE
Family Name:	Copy of 05-06vdisk1
World Wide LUN Name:	6005-08b4-0001-0a70-0002-0000
UUID:	6005-08b4-0001-0a70-0002-0000-0048-0000
Condition/State	
Operational State:	Good
Date/Time	
Created:	14-Aug-2003 16:31:34
Cache Policies	
Write:	Mirrored write-back
Read:	On
Attributes	
Type:	Original
Disk Group:	Disk Group 2
Capacity Req:	1 GB
Capacity Used:	1 GB
Redundancy:	Vraid1
Write Protect:	No

Figure 75: Changing Write Protect attribute of Vdisk

- b. Click the drop-down arrow for the **Write Protect** attribute and select **No**.
- c. Choose the **Save changes** tab. An Operations succeeded page is displayed. Click **OK**.

9. On the Vdisk Active Member Properties page, either change the World Wide LUN ID Name in the box in the Identification pane (Figure 75) or use the default name that is provided.

For example, if you are using an archival application that requires the presentation of the Snapcloned Vdisks to have the same World Wide LUN ID names, then change the LUN ID to a name of your choice, however it must start with the number 6. All remaining numbers can be hex digits. (For example, 6234-1234-1234-1234-1234-1234-1234-1234.) This World Wide LUN ID would then be used during all subsequent backup operations using the moved Vdisks. After entering the new LUN ID, choose the **Save changes** tab. A message is displayed asking for confirmation of the name change. Click **OK**.

If you accept the default World Wide LUN ID name, proceed to the next step.

10. An Operations succeeded page is displayed upon completion of the LUN ID name change. Click **OK**.
11. Present the moved Vdisks on the remote storage system (HSV18) to hosts.
 - a. Select the active member of a moved Vdisk. The Vdisk Active Member Properties page is displayed (Figure 75).
 - b. Choose the **Presentation** tab. The Vdisk Active Member Properties page is displayed with a **Present** tab (Figure 76).

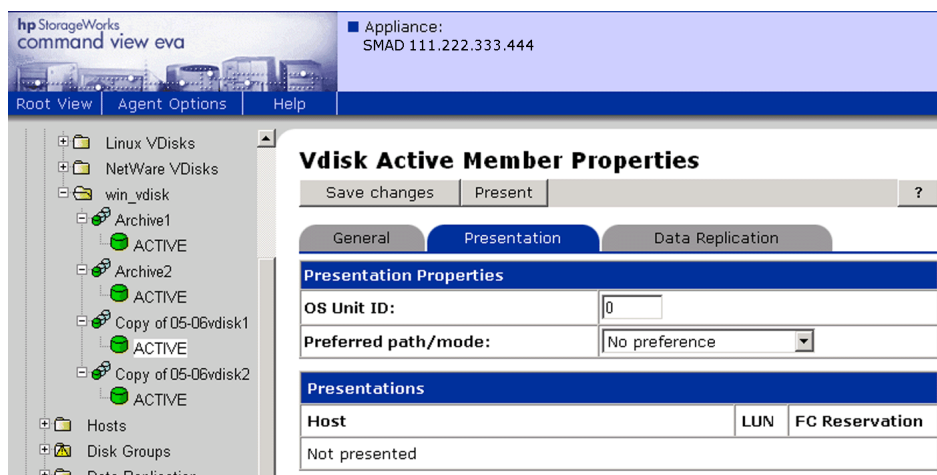


Figure 76: Vdisk Active Member Properties page with Present tab

- c. Click the **Present** tab. The Present Vdisk page is displayed (Figure 77).

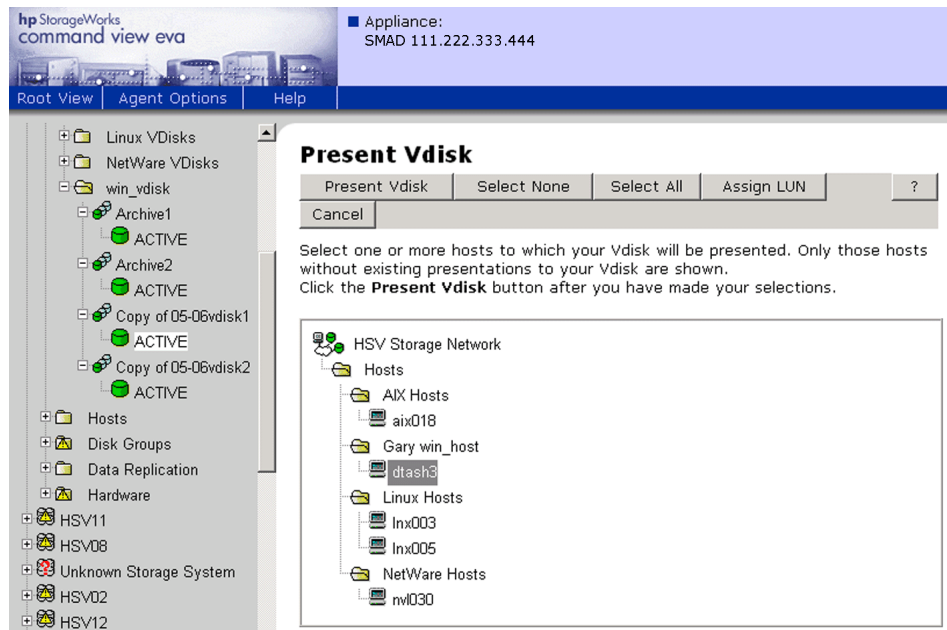


Figure 77: Present Vdisk page

- d. Select an available host, then choose the **Present Vdisk** tab. An Operation succeeded page is displayed.
- e. Click **OK**.
12. Rename the Vdisk to a useful name.
- Select the Vdisk in the navigation pane. The Vdisk Family Properties page is displayed (Figure 78).
 - Type a new name in the Name field and choose the **Save changes** tab. An Operation succeeded page is displayed.
 - Click **OK**.




Figure 78: Vdisk Family Properties page

The Vdisks are now available on the remote storage system for any purpose.

Bootless DR group planned failover with Linux using LVM in standalone mode or with SuSE SLES 8 running LifeKeeper 4.4.3

The following procedures describe how to perform a bootless DR group failover when running the Logical Volume Manager (LVM) with Linux. Separate procedures for running in standalone host mode or with clusters (LifeKeeper) are listed. Perform the procedures for the source host, followed by the procedures for the destination host.

Note: This procedure is not supported for unplanned failovers. The term “bootless” means that after the LUNs are first presented to a destination host, which requires an initial reboot, no further reboot of that host should be required.

Source host procedure

Perform one of the following steps on the source host, depending on whether or not you are running LifeKeeper 4.4.3.

1. If you are running LifeKeeper 4.4.3, proceed to step 2. If you are not running LifeKeeper, perform the following steps:
 - a. From your source host, stop I/O to your LUNs. Allow enough time for the I/O to complete before proceeding to the next step.
 - b. Unmount the volumes contained in the DR group.
Example: `umount /mounts/lvol1`
 - c. Change the status of the LUNs to inactive with the following command:
`vgchange VolumeGroupName -a n`
Example: `vgchange vg01 -a n`
 - d. Make the group unknown to the system with the `vgexport` command.
Example: `vgexport vg01`
 - e. Perform a failover of the DR group using the Continuous Access user interface.

- f. Depending on the number of LUNs, do one of the following to prevent Secure Path from detecting a failed disk:
 - For individual LUNs, run `spmgr quiesce -p path` (for each path visible to the LUNs).
 - For all LUNs at once, run `spmgr set -p off`. This method will turn off path verification for all LUNs still visible to the system.
2. If you are running LifeKeeper 4.4.3 clusters:
 - a. Bring your resources “out of service” with the LifeKeeper GUI.
 - b. From the system console:
 - 1) Enter the `mount` command to verify the volume is unmounted.
 - 2) Enter the `vgscan` command to verify that the volume group was exported.
 - c. Perform a failover of the DR group with the Continuous Access user interface.
 - d. Depending on the number of LUNs, do one of the following to prevent Secure Path from detecting a failed disk:
 - For individual LUNs, run `spmgr quiesce -p path` (for each path visible to the LUNs).
 - For all LUNs at once, run `spmgr set -p off`. This method will turn off path verification for all LUNs still visible to the system.

Destination host procedure

If this is the first time that LUNs are being presented to the destination host, reboot the host to pick up the new LUNs. If a reboot is not required (LUNs have been previously presented), and the paths are quiesced, issue the `spmgr restart all` command to unquiesce the paths.

Perform one of the following steps on the destination host, depending on whether or not you are running LifeKeeper 4.4.3.

1. If you are running LifeKeeper 4.4.3, proceed to step 2. If you are not running LifeKeeper, perform the following steps:
 - a. Issue the following command to make the volume known to the system:

```
vgimport VolumeGroupName PhysicalVolumePath
```

Example: `vgimport vg01 /dev/sda1`

- b. Mount the file systems.

Example: `mount -t reiserfs /dev/vg01/lvol1 /mounts/lvol1`

- c. Start host I/O.
- d. If the verification path is turned off, issue the following command:

`spmgr set -p on`

- 2. If you are running LifeKeeper 4.4.3 clusters:

- a. If this is the first time LUNs are being presented to the destination host, you must build the resource hierarchies for each new LUN presented (refer to the documentation on the LifeKeeper CD).
- b. Bring your resources “out of service” with the LifeKeeper GUI.
- c. Start host I/O.
- d. If the verification path is turned off, issue the following command:

`spmgr set -p on`

Performing a failover back to the previous source

After performing a failover from a source to a destination storage system, a failover can be performed in the other direction after waiting a minimum of 15 minutes. Use the same procedures described in the previous paragraphs for the source and destination storage systems.

Description of Event and Termination Codes



This appendix describes event and termination codes that you may encounter.

Viewing event codes may be helpful for troubleshooting purposes. Events generated by the Enterprise Virtual Array are sent to the Storage Management Appliance. These events can then be viewed in Command View EVA by selecting a storage system, choosing **View Events**, then clicking **Controller Event Log**. [Table 11](#) lists events applicable to Continuous Access EVA.

If your storage system fails, a termination code is generated. Termination codes can be viewed in Command View EVA by selecting a storage system, choosing **View Events**, then clicking **Termination Event Log**. [Table 12](#) lists termination events applicable to Continuous Access EVA.

Command View EVA can be configured to send events applicable to Continuous Access EVA to the Continuous Access user interface. (Refer to the *HP StorageWorks Continuous Access EVA Installation Guide*.) These events are displayed in summary format in the Continuous Access user interface **View Events** tab. The event codes are the same as in Command View EVA. The event code display is for convenience only, and any troubleshooting should be done from the Command View EVA Controller Event Log. [Table 13](#) lists event codes applicable to Continuous Access EVA and a summary of the event as described by the Continuous Access user interface.

Table 11: Continuous Access EVA event codes with Command View EVA descriptions

Event code	Command View EVA description
0946000e	Severity: Normal -- informational in nature. The Data Replication Group identified in the handle field was created.
0947000e	Severity: Normal -- informational in nature. The Data Replication Group identified in the handle field was deleted.
0971000f	<p>Severity: Normal -- informational in nature. The HSV110 controller identified in the handle field received a request to shutdown. <P> The old_attr.type field and the new_attr.type fields contain the value 1. The old_attr.value.u32[0] field indicates the type of restart that was requested as follows:</p> <ul style="list-style-type: none"> 0 = None -- no restart 1 = Regular -- full restart, host system connectivity is lost until the controller returns to normal operation 2 = Fast -- resynchronization, restart of the controller in a manner that has little or no impact on host system connectivity <P> The old_attr.value.u32[1] field indicates whether the other HSV110 controller of the pair was requested to remain operational or to also shutdown as follows: 0 = Remain operational 1 = Coupled shutdown <P> The old_attr.value.u32[2] field indicates whether the HSV110 controller was requested to remain in the power on state or power itself off as follows: 0 = Remain in the power on state 1 = Power itself off <P> The old_attr.value.u32[3] field value indicates whether the Physical Disk Drive enclosure(s) was requested to remain in the power on state or to be powered off as follows: 0 = Remain in the power on state 1 = Powered off <P> The old_attr.value.u32[4] field value indicates whether the battery assemblies located within the HSV110 controller were requested to be enabled or disabled as follows: 0 = Enabled 1 = Disabled <P> The old_attr.value.u32[5] field contains the number of seconds the shutdown operation was requested to be delayed.

Table 11: Continuous Access EVA event codes with Command View EVA descriptions (Continued)

Event code	Command View EVA description
0972000f	<p>Severity: Normal -- informational in nature. The HSV110 controller identified in the handle field has completed its shutdown preparations. <P> The old_attr.type field and the new_attr.type fields contain the value 1. The old_attr.value.u32[0] field value indicates the result of the HSV110 controller Cache Memory shutdown operation as follows:</p> <ul style="list-style-type: none"> 1 = Success 2 = Failure 3 = Not applicable <p><P> The old_attr.value.u32[1] field contains the internal Storage System Management Interface cache shutdown status code. <P> The old_attr.value.u32[2] field value indicates the result of the Physical Disk Drive enclosure(s) power off operation as follows:</p> <ul style="list-style-type: none"> 1 = Success 2 = Failure 3 = Not applicable <p><P> The old_attr.value.u32[3] field contains the internal Storage System Management Interface Physical Disk Drive enclosure(s) power off status code. <P> The old_attr.value.u32[4] field value indicates the result of the battery assemblies disable operation as follows:</p> <ul style="list-style-type: none"> 1 = Success 2 = Failure 3 = Not applicable <p><P> The old_attr.value.u32[5] field value indicates the battery assemblies disable operation failure mode as follows: </p> <ul style="list-style-type: none"> 0 = No failure indicated. 1 = Failed only on the HSV110 controller identified in the handle field. 2 = Failed only on the other HSV110 controller of the pair. 3 = Failed on both HSV110 controllers.

Table 11: Continuous Access EVA event codes with Command View EVA descriptions (Continued)

Event code	Command View EVA description
0973000f	<p>Severity: Normal -- informational in nature. The Failsafe state of the Data Replication Group identified in the handle field has changed. The old_attr.value.u32[0] field contains the old Failsafe state. The new_attr.value.u32[0] field contains the new Failsafe state. The Failsafe state values that may be found in the old_attr.value.u32[0] and new_attr.value.u32[0] fields are as follows:</p> <ul style="list-style-type: none"> 1 = Failsafe Disabled 2 = Failsafe Enabled
0974000f	<p>Severity: Normal -- informational in nature. The Mode of the Data Replication Group identified in the handle field has changed. The old_attr.value.u32[0] field contains the old Mode. The new_attr.value.u32[0] field contains the new Mode. The Mode values that may be found in the old_attr.value.u32[0] and new_attr.value.u32[0] fields are as follows:</p> <ul style="list-style-type: none"> 0 = Normal Active Source 1 = Normal Active Destination 2 = Active/Active (Master) 3 = Active/Active (Non-Master)
0975000f	<p>Severity: Normal -- informational in nature. The Operational state of a Data Replication Group has changed to Synchronous or Asynchronous. The old_attr.value.u32[0] field contains the old Operation state. The new_attr.value.u32[0] field contains the new Operation state. The Operation state values that may be found in the old_attr.value.u32[0] and new_attr.value.u32[0] fields are as follows:</p> <ul style="list-style-type: none"> 1 = Synchronous 2 = Asynchronous

Table 11: Continuous Access EVA event codes with Command View EVA descriptions (Continued)

Event code	Command View EVA description
0976000f	<p>Severity: Normal -- informational in nature. The Read Only attribute of the Data Replication Group identified in the handle field has changed. The old_attr.value.u32[0] field contains the old Read Only attribute. The new_attr.value.u32[0] field contains the new Read Only attribute. The Read Only attribute values that may be found in the old_attr.value.u32[0] and new_attr.value.u32[0] fields are as follows:</p> <ul style="list-style-type: none"> 0 = Data Replication Destination Storage System Virtual Disk disabled for read access. 1 = Data Replication Destination Storage System Virtual Disk enabled for read access.
0977000f	<p>Severity: Normal -- informational in nature. A Data Replication Group failover has occurred. The handle field contains the identity of the Data Replication Group. The old_attr.value.u32[0] field contains the old Role. The new_attr.value.u32[0] field contains the new Role.</p> <ul style="list-style-type: none"> 0 = Normal Active Source 1 = Normal Active Destination 2 = Active/Active (Master) 3 = Active/Active (Non-Master)
0978000f	<p>Severity: Normal -- informational in nature. A Data Replication Group has been suspended or resumed. The handle field contains the identity of the Data Replication Group. The old_attr.value.u32[0] field contains the old Suspend state. The new_attr.value.u32[0] field contains the new Suspend state. The Suspend state values that may be found in the old_attr.value.u32[0] and new_attr.value.u32[0] fields are as follows:</p> <ul style="list-style-type: none"> 1 = Connection between the Data Replication Source and Data Replication Destination is active. 2 = Connection between the Data Replication Source and Data Replication Destination is inactive.
0979000f	<p>Severity: Normal -- informational in nature. A Storage System Virtual Disk was added to a Data Replication Group. The Storage System Virtual Disk is identified in the add_handle field. The Data Replication Group is identified in the handle field.</p>

Table 11: Continuous Access EVA event codes with Command View EVA descriptions (Continued)

Event code	Command View EVA description
097a000f	Severity: Normal -- informational in nature. A Storage System Virtual Disk was removed from a Data Replication Group. The Storage System Virtual Disk is identified in the add_handle field. The Data Replication Group is identified in the handle field.
09c85105	Severity: Undetermined -- more information needed to determine severity. The Data Availability state of the internal Logical Disk associated with the Virtual Disk identified in the handle field has transitioned to the DATA LOST state. The value.ul1 field contains the new state: 1 (DATA LOST). The value.ul2 field contains the old state: 0 (NORMAL).
09c95105	Severity: Undetermined -- more information needed to determine severity. The state of the Disk Group identified in the handle field has transitioned to an INOPERATIVE state. The value.ul1 field contains the new Disk Group state. The value.ul2 field contains the old Disk Group state. The state values that may be found in the value.ul1 and value.ul2 fields are as follows: 1 = Normal 2 = Disk Group with no redundancy is inoperative 3 = Disk Group with parity redundancy is inoperative 4 = Disk Group with mirrored redundancy is inoperative 5 = Disk Group with no redundancy is inoperative, marked for re-use 6 = Disk Group with parity redundancy is inoperative, marked for re-use 7 = Disk Group with mirrored redundancy is inoperative, marked for re-use

Table 11: Continuous Access EVA event codes with Command View EVA descriptions (Continued)

Event code	Command View EVA description
09ca5105	<p>Severity: Undetermined -- more information needed to determine severity. The state of the internal Logical Disk associated with the Virtual Disk identified in the handle field has transitioned to the FAILED state. The value.ul1 field contains the new state: 5 (FAILED). The value.ul2 field contains the old state. The state values that may be found in the value.ul2 field are as follows: </p> <ul style="list-style-type: none"> 1 = Normal 2 = Replacement delay in progress 3 = Redundancy lost, restore in progress 4 = Redundancy lost, restore stalled 6 = Creation in progress 7 = Snapshot is inoperative due to lack of snapshot space 8 = Deletion in progress 9 = Capacity change in progress 10 = Inoperative due to data lost 11 = Capacity reservation in progress 12 = Capacity unreservation in progress <p></p>
09cc5105	<p>Severity: Undetermined -- more information needed to determine severity. The state of the internal Logical Disk associated with the Virtual Disk identified in the handle field has transitioned to the DEVICE DATA LOST state. The value.ul1 field contains the new state: 10 (DEVICE DATA LOST). The value.ul2 field contains the old state. The state values that may be found in the value.ul2 field are as follows: </p> <ul style="list-style-type: none"> 1 = Normal 2 = Replacement delay in progress 3 = Redundancy lost restore, in progress 4 = Redundancy lost restore, stalled 5 = Failed 6 = Creation in progress 7 = Snapshot is inoperative due to lack of snapshot space 8 = Deletion in progress 9 = Capacity change in progress 11 = Capacity reservation in progress 12 = Capacity unreservation in progress <p></p>

Table 11: Continuous Access EVA event codes with Command View EVA descriptions (Continued)

Event code	Command View EVA description
09cdc305	Severity: Undetermined -- more information needed to determine severity. The state of the Fibre Channel port identified in the attribute.value.str field and located on the rear panel of the HSV110 controller identified in the handle field has transitioned to the FAILED state. The attribute.type field contains the value 4. The value.ul1 field contains the new state: 2 (FAILED). The value.ul2 field contains the old state: 1 (NORMAL). <P> Note that the HSV110 controller's internal identity of the affected Fibre Channel port is contained in the secondary_id field.
09ce0005	Severity: Normal -- informational in nature. The state of the Disk Group identified in the handle field has transitioned to an INOPERATIVE MARKED state. The value.ul1 field contains the new Disk Group state. The value.ul2 field contains the old Disk Group state. The state values that may be found in the value.ul1 and value.ul2 fields are as follows: 1 = Normal 2 = Disk Group with no redundancy is inoperative 3 = Disk Group with parity redundancy is inoperative 4 = Disk Group with mirrored redundancy is inoperative 5 = Disk Group with no redundancy is inoperative, marked for re-use 6 = Disk Group with parity redundancy is inoperative, marked for re-use 7 = Disk Group with mirrored redundancy is inoperative, marked for re-use
09cf4105	Severity: Warning -- not failed but attention recommended or required. The state of the Physical Disk Drive identified in the handle field has transitioned to the NOT PRESENT state. The value.ul1 field contains the new state: 4 (NOT PRESENT). The value.ul2 field contains the old state. The state values that may be in the value.ul2 field are as follows: 1 = Normal 2 = Degraded 3 = Failed 5 = Single Port on Fibre

Table 11: Continuous Access EVA event codes with Command View EVA descriptions (Continued)

Event code	Command View EVA description
09d35105	Severity: Undetermined -- more information needed to determine severity. At least one Virtual Disk associated with the Data Replication Group identified in the handle field has transitioned to the INOPERATIVE state. The remaining Virtual Disks associated with this Data Replication Group have been forced INOPERATIVE. The value.ul1 field contains the new Data Replication Group member state: 2 (INOPERATIVE). The value.ul2 field contains the old Data Replication Group member state: 1 (OPERATIVE).
09d40005	Severity: Normal -- informational in nature. All the Virtual Disks associated with the Data Replication Group identified in the handle field have transitioned to the OPERATIVE state. The value.ul1 field contains the new Data Replication Group member state: 1 (OPERATIVE). The value.ul2 field contains the old Data Replication Group member state: 2 (INOPERATIVE).
09d50005	Severity: Normal -- informational in nature. The state of the Physical Disk Drive identified in the handle field has transitioned to the Single Port on Fibre state. The value.ul1 field contains the new state: 5 (SINGLE PORT ON FIBRE). The value.ul2 field contains the old state. The state values that may be in the value.ul2 field are as follows: 1 = Normal 2 = Degraded 3 = Failed 4 = Not Present
0c03000c	Severity: Normal -- informational in nature. The specified Data Replication Group has transitioned to the Merging state, because the alternate Storage System or Destination Virtual Disk is now accessible or resumed. The following fields are not valid: status, blocks, vda.
0c045f0c	Severity: Critical -- failure or failure imminent. A Data Replication Group has entered the Failsafe Locked state because the Data Replication Destination Site is inaccessible. The following fields are not valid: status, blocks, vda, port, cerp_id, side.
0c05600c	Severity: Critical -- failure or failure imminent. A Data Replication Group has entered the Failsafe Locked state due to an inoperative source. The following fields are not valid: status, blocks, vda, port, cerp_id, side.

Table 11: Continuous Access EVA event codes with Command View EVA descriptions (Continued)

Event code	Command View EVA description
0c06600c	Severity: Critical -- failure or failure imminent. A Full Copy was terminated prior to completion. An unrecoverable read error occurred on the specified Source Virtual Disk during the Full Copy.
0c075f0c	Severity: Critical -- failure or failure imminent. A Full Copy terminated prior to completion. A remote copy error occurred due to an inaccessible alternate Storage System. The Full Copy will continue when the Data Replication Destination is restored.
0c08610c	Severity: Critical -- failure or failure imminent. A Full Copy terminated prior to completion. A remote copy error occurred due to an inaccessible Destination Virtual Disk. The Full Copy will continue when the Data Replication Destination is restored.
0c09620c	Severity: Warning -- not failed but attention recommended or required. A Data Replication Log has been reset due to insufficient Disk Group capacity. The Data Replication Destination has been marked for a Full Copy. The following fields are not valid: status, blocks, vda, port, cerp_id, side.
0c0a000c	Severity: Normal -- informational in nature. A Data Replication Log has been reset due to a Data Replication Group failover. The following fields are not valid: status, blocks, vda.
0c0c000c	Severity: Normal -- informational in nature. A Destination Data Replication Group has successfully completed a Merge. The following fields are not valid: status, blocks, vda.
0c0f000c	Severity: Normal -- informational in nature. A Data Replication Group is no longer in a Failsafe Locked state. The following fields are not valid: status, blocks, vda, port, cerp_id, side.
0c10000c	Severity: Normal -- informational in nature. A Destination Data Replication Group has been marked for a Full Copy. The following fields are not valid: status, blocks, vda.
0c11000c	Severity: Normal -- informational in nature. A Storage System has discovered that a Data Replication Group failover has taken place. This Virtual Disk is transitioning from a Data Replication Source role to a Data Replication Destination role. The following fields are not valid: status, blocks, vda.
0c12000c	Severity: Normal -- informational in nature. This Data Replication Group is transitioning from a Data Replication Destination role to a Data Replication Source role. The following fields are not valid: status, blocks, vda.

Table 11: Continuous Access EVA event codes with Command View EVA descriptions (Continued)

Event code	Command View EVA description
0c155f0c	Severity: Critical -- failure or failure imminent. The Data Replication Path between this Site and the Alternate Site has closed, possibly due to a connection failure between the specified host port and the Alternate Site. The following fields are not valid: group_name_uuid, group_uuid, source_scvd_uuid, status, block, vda.
0c17630c	Severity: Critical -- failure or failure imminent. The Data Replication Manager communications protocol version between the Data Replication Source <REFERENCE>(DRM_SITE) and a Data Replication Destination <REFERENCE>(DRM_SITE) is mismatched.
0c18640c	Severity: Critical -- failure or failure imminent. Conditions on the Data Replication Destination Site are preventing acceptable replication throughput. Initiating temporary logging on the affected Data Replication Group that is failsafe mode disabled.
0c1a000c	Severity: Normal -- informational in nature. The specified Destination Virtual Disk has successfully completed a Full Copy. The following fields are not valid: status, blocks, vda.
0c1b5f0c	Severity: Critical -- failure or failure imminent. A Data Replication Group has transitioned to the Logging state because the alternate Storage System is not accessible.
0c1c610c	Severity: Critical -- failure or failure imminent. The specified Source Data Replication Group has transitioned to the Logging state because the Destination Virtual Disk is not accessible. The following fields are not valid: status, blocks, vda, port, cerp_id, side.
0c1d000c	Severity: Normal -- informational in nature. A software problem was found in the group log. A Full Copy of the affected Data Replication Group will be initiated.
0c1e5f0c	Severity: Critical -- failure or failure imminent. The members of the specified Source Data Replication Group have not been presented to the host because the remote Storage System is not accessible. Suspend Source Data Replication Group to override this behavior, which will present the members.

Table 11: Continuous Access EVA event codes with Command View EVA descriptions (Continued)

Event code	Command View EVA description
0c1f000c	Severity: Normal -- informational in nature. The members of the specified Source Data Replication Group have been presented to the host because the remote Storage System is now accessible, or source group is now suspended.
0c20650c	Severity: Critical -- failure or failure imminent. Conditions on the Data Replication Destination Site are preventing replication processing. The specified Source Data Replication Group will remain in the Logging or the Failsafe Locked state until corrective action is performed.
0c21660c	Severity: Critical -- failure or failure imminent. Conditions on the Data Replication Source Site are preventing replication processing. The specified Source Data Replication Group will remain in the Logging or the Failsafe Locked state until corrective action is performed.

Table 12: Continuous Access EVA termination codes and descriptions

Termination code	Description
0c010100	Severity: Critical -- failure or failure imminent. Invalid Data Replication Manager Dual State was used for MFC communication between the dual controllers.
0c030100	Severity: Critical -- failure or failure imminent. Invalid state exists for deleting a Group State Block.
0c040100	Severity: Critical -- failure or failure imminent. A software problem was found in processing a recovery write upon controller start or failover. The group sequence number node already exists.
0c050100	Severity: Critical -- failure or failure imminent. A software problem was found in processing a recovery write upon controller start or failover. The recovery write data was not in cache as expected.
0c060100	Severity: Critical -- failure or failure imminent. A software problem was found in processing a recovery write upon controller start or failover. The recovery write data found in cache was not marked dirty write-back cached data as expected.
0c070100	Severity: Critical -- failure or failure imminent. A software problem was found in processing a recovery write upon controller start or failover. Lookup of group sequence number node failed.
0c080100	Severity: Critical -- failure or failure imminent. A software problem was found in cleaning Data Replication Manager context in the mirror cache when the mirror was declared invalid. A recovery write was found, but its associated RIE was not marked free as expected.
0c090100	Severity: Critical -- failure or failure imminent. A software problem was found in cleaning Data Replication Manager context in the mirror cache when the mirror was declared invalid. Not all group members were processed.
0c0a0100	Severity: Critical -- failure or failure imminent. A software problem was found in cleaning Data Replication Manager context in the primary cache when the primary was declared invalid. A recovery write was found, but its associated RIE was not marked free as expected.
0c0b0100	Severity: Critical -- failure or failure imminent. A software problem was found in cleaning Data Replication Manager context in the primary cache when the primary was declared invalid. Not all group members were processed.
0c0c0100	Severity: Critical -- failure or failure imminent. A software problem was found when deleting the Group State Block. Transfers were not completely run down.

Table 12: Continuous Access EVA termination codes and descriptions (Continued)

Termination code	Description
0c0d0100	Severity: Critical -- failure or failure imminent. A software problem was found when inserting a Group State Block into the active list. A Group State Block with this same Universal Unique Identifier is already on the active list.
0c0e0100	Severity: Critical -- failure or failure imminent. A group sequence number out of order was detected in the transfer path upon remote write completion after the mirror controller was updated. A Full Copy of the affected Data Replication Group may be initiated upon the next controller restart.
0c0f0100	Severity: Critical -- failure or failure imminent. Setting the e-bit failed for a write long command on the destination unit.
0c100100	Severity: Critical -- failure or failure imminent. An attempt was made to acquire the Data Replication Manager Remote Response Waiter, but it was unexpectedly already in use.
0c110100	Severity: Critical -- failure or failure imminent. A Group Sequence Number Node was lost during mirror synchronization.
0c140100	Severity: Critical -- failure or failure imminent. A group sequence number out of order was detected in the transfer path on the mirror side upon remote write completion. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.
0c150100	Severity: Critical -- failure or failure imminent. A group sequence number out of order was detected upon controller restart or failover when building the list of incomplete writes. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.
0c160100	Severity: Critical -- failure or failure imminent. A group sequence number out of order was detected upon controller restart or failover when completing previously incomplete writes. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.
0c170100	Severity: Critical -- failure or failure imminent. A group sequence number out of order with the received group sequence number was detected after a controller restarted, when synchronizing the group sequence numbers with the mirror side. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.

Table 12: Continuous Access EVA termination codes and descriptions (Continued)

Termination code	Description
0c180100	Severity: Critical -- failure or failure imminent. A group sequence number out of order with the use and sent group sequence numbers was detected after a controller restarted, when synchronizing the group sequence numbers with the mirror side. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.
0c190100	Severity: Critical -- failure or failure imminent. A group sequence number out of order with the received group sequence number was detected after a controller restarted, when synchronizing the group sequence numbers with the primary side. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.
0c1a0100	Severity: Critical -- failure or failure imminent. A group sequence number out of order with the use and sent group sequence number was detected after a controller restart, when synchronizing the group sequence numbers with the primary side. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.
0c1b0100	Severity: Critical -- failure or failure imminent. A group sequence number out of order with the use group sequence number was detected after a controller restart, when synchronizing the group sequence numbers with the primary side. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.
0c1c0100	Severity: Critical -- failure or failure imminent. A group sequence number out of order was detected after a controller restart when synchronizing the mirror writes with the primary side. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.
0c200100	Severity: Critical -- failure or failure imminent. A group sequence number out of order with the use group sequence number was detected after a controller restart, when synchronizing the group sequence numbers with the primary side. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.
0c210100	Severity: Critical -- failure or failure imminent. A group sequence number out of order with the use group sequence number was detected to be too high. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.

Table 12: Continuous Access EVA termination codes and descriptions (Continued)

Termination code	Description
0c220100	Severity: Critical -- failure or failure imminent. A group sequence number out of order with the use group sequence number was detected to be too low. A Full Copy of the affected Data Replication Group may be initiated upon controller restart.
0c230100	Severity: Critical -- failure or failure imminent. A Data Replication Group member was detected to be in an unexpected cache state.
0c240100	Severity: Critical -- failure or failure imminent. Secondary controller selection failed

Table 13: Event codes with Continuous Access user interface summary description

Event code	Continuous Access user interface description
0946000e	A DR group was created.
0947000e	A DR group was deleted.
0971000f	A controller has received a request to shutdown.
0972000f	A controller has completed its shutdown preparations.
0973000f	The failsafe state of a DR group has been enabled or disabled.
0974000f	The replication mode of a DR group has changed at the source or destination.
0975000f	The operational state of a DR group has changed to synchronous or asynchronous.
0976000f	Access to a destination Vdisk has been enabled or disabled.
0977000f	A DR group failover has occurred.
0978000f	A DR group has been suspended or resumed.
0979000f	A Vdisk was added to a DR group.
097a000f	A Vdisk was removed from a DR group.
09c85105	A Vdisk has lost data.
09c95105	A Disk group has gone inoperative.
09ca5105	A Vdisk has failed.
09cc5105	A Vdisk has lost data.
09cdc305	A Fibre Channel port located on the rear panel of the controller has failed.
09ce0005	A Disk group has gone inoperative.
09cf4105	A physical disk drive is no longer present.
09d35105	A DR group member has gone inoperative, and the remaining members have been forced inoperative.
09d40005	A DR group has become operative.
09d50005	A physical disk drive is only accessible on a single port.
0c00000c	A destination Vdisk has successfully completed a full copy.
0c015f0c	A DR group is logging.
0c02610c	A DR group is logging because the destination Vdisk is not accessible.

Table 13: Event codes with Continuous Access user interface summary description (Continued)

Event code	Continuous Access user interface description
0c03000c	A source Vdisk is now merging, because the alternate storage system or destination Vdisk is now accessible or resumed.
0c045f0c	A DR group has entered the failsafe locked state.
0c05600c	A DR group has entered the failsafe locked state.
0c06600c	An unrecoverable read error occurred during a full copy. The full copy was terminated.
0c075f0c	A full copy terminated prior to completion. A remote copy error occurred due to an inaccessible destination storage system. The full copy will continue when the destination storage system is restored.
0c08610c	A full copy terminated prior to completion. A remote copy error occurred due to an inaccessible destination storage system. The full copy will continue when the destination Vdisk is restored.
0c09620c	A log has been reset due to insufficient disk group capacity. The destination Vdisk has been marked for a full copy.
0c0a000c	A log has been reset due to a DR group failover.
0c0c000c	A destination Vdisk has successfully completed a merge.
0c0f000c	A DR group is no longer in a failsafe locked state.
0c10000c	A destination Vdisk has been marked for a full copy.
0c11000c	A storage system has just discovered that a DR group failover has taken place. A Vdisk is transitioning from a source role to a destination role.
0c12000c	A DR group member is transitioning from a destination role to a source role.
0c15000c	A data replication path between this site and the alternate site has closed.

Glossary

This glossary defines terms used in this guide or related to this product and is not a comprehensive glossary of computer terms.

actual disk failure protection level

The actual level of protection available to the disk group as specified in its current configuration. If the actual failure protection level is less than the level you need, add more physical disks to the disk groups.

allocation policy

Storage system rules that govern how virtual disks are created. There are two rules:

- Allocate completely—The space a virtual disk requires on the physical disks is reserved, even if the virtual disk is not currently using the space.
- Allocate on demand—The space a virtual disk requires on the physical disks is not reserved until needed.

alternate site

Controller pair or storage system containing a copy of the data stored at a primary site. A storage system may be considered a primary site for some LUNs and an alternate site for other LUNs. An alternate site could be in the same room or building as the primary site, but generally it is not.

array

All the physical disk drives in a storage system that are known to and under the control of a controller pair. Multiple arrays can exist in the same cabinet, or one array can include more than one cabinet (as in the case of an expansion cabinet).

asynchronous replication mode

The mode of operation of a DR group whereby the EVA controller firmware provides I/O completion acknowledgement to the host after data is delivered in cache at the source, and before the data delivery to cache on the destination. This mode is not available at this time. *See also* DR Group write mode.

B-series switches

Fibre Channel core and SAN switches made by Brocade and sold by HP. Refer to the *HP StorageWorks Continuous Access EVA Design Reference Guide* for a listing of these switches.

bandwidth

The transmission capacity of a link or system, usually measured in bits per second.

bidirectional

The circumstance when a storage system is configured so that it contains both source and destination Vdisks. This configuration allows for multidirectional I/O flow among several storage systems.

C-series switches

Fibre Channel switches made by Cisco and sold by HP. Refer to the *HP StorageWorks Continuous Access EVA Design Reference Guide* for a listing of these switches.

CAC

Corrective action code. A display component that defines the action required to correct a problem. This component is displayed on the Command View EVA graphical user interface (GUI).

Command View EVA

Command View EVA consists of:

- The graphical user interface (GUI) that displays the usage of the storage system.
- The software behind the GUI, which controls and monitors the functions.

The Command View EVA software can be installed on more than one Storage Management Appliance in a fabric. Each installation of the Command View EVA software is a management agent. The client for the agent is a standard browser.

Continuous Access EVA

Continuous Access EVA is a storage-based HP StorageWorks product consisting of two or more storage systems performing disk-to-disk replication, along with the management user interfaces that facilitates configuring, monitoring, and maintaining the replicating capabilities of the storage systems.

Continuous Access user interface

A tool for managing the replication of storage objects in a SAN. It provides a graphical user interface for the management of disk I/O, failover, and maintenance operations. The Continuous Access user interface is installed in the HP OpenView Storage Management Appliance and appears as a software service.

copy set

A generic term for a logical disk in one storage array that is replicated to another logical disk in another storage array. There are two states: *normal* and *copying*. The term is commonly used to represent the pair of Vdisks, one on the source array and one on the destination array.

corrective action code

See CAC.

data distribution

Pushing copies of data to geographic locations to make it more easily accessible to many customers.

data entry mode

The mode in which controller information can be displayed or controller configuration data can be entered. On the Enterprise storage system, the data entry mode is active when the LCD on the HSV controller operator control panel is flashing.

data migration

Moving data to a new location or to one with a larger capacity.

data movement

Continuous Access EVA provides for data movement services, such as data backup, data migration, data distribution, and data mining.

data replication mode

See DR mode.

default disk group

The first disk group created when the storage system is initialized. The default disk group can contain up to the entire set of physical disks in the array. The minimum number of physical disks the default disk group can contain is eight. The maximum is the number of installed disks and is limited to 14 drives per attached drive enclosure.

destination Vdisk

A Vdisk that is the recipient of replicated data from a source Vdisk.

destination Vdisk access

A storage system's ability to allow host access to a destination Vdisk. There are two options:

- Disabled
- Read-only (write-protected)

disaster tolerance

The capability for rapid recovery of user data from a remote location when a significant event or disaster occurs at the local computing site. It is a special combination of high-availability technology and services that can continue the operation of critical applications in the event of a site disaster. Disaster-tolerant systems are designed to allow applications to continue operating during the disaster recovery period.

disk failure protection

The three levels of disk failure protection are:

- **None**—No protection is present.
- **Single**—The capacity of one physical disk is reserved.
- **Double**—The capacity of two physical disks is reserved.

Disk failure protection occurs when the storage system sets aside reserved capacity to take over the functionality of a failed or failing physical disk drive. In groups with mixed capacity drives, the reserved capacity is based on the largest disk in the disk group. The system must cover a failure in any drive, so it reserves enough capacity to cover the largest failure that could happen.

disk group

A named group of physical disks selected from all the available physical disks in the storage system where one or more Vdisks can be created. A physical disk can belong to only one disk group.

disk group occupancy alarm level

An event code that is generated when the amount of data stored in the disk group reaches a peak level of the total disk group capacity. For example, if a disk group's capacity is 288 GB, and the occupancy alarm level is 80%, an event code is generated when the amount of data in the disk group reaches 230.4 GB. The default occupancy alarm level is 95% of the total disk group capacity.

disk migration state

A physical disk drive operating state. Two states are possible:

- **Stable**—The physical disk drive has not failed.
- **Migrating**—The disk drive is failing or failure is predicted. If this state occurs, data moves from the disk onto other disk drives in the same disk group.

disk replacement delay

The time that elapses between detection of a possible drive failure and when the controller starts searching for spare disk space. Drive replacement seldom starts immediately, in case the failure was a temporary condition.

distributed sparing

Allocated space per disk group to recover from physical disk failures in that disk group.

DR group

A VCS construct organizing one or more Vdisks in an HSV storage system so that they replicate to the same specified destination, fail over together if a single Vdisk in the collection fails, and preserve write-ordering within the collection. This is the HSV implementation of an association set, and its member Vdisks are the HSV implementation of copy sets.

DR group direction

The replication direction of a DR group. There are two states:

- Original (from *Home*)
- Reversed (failed over, or toward *Home*)

DR group log state

The current behavior of the log associated with a DR group. In the state options, references to multiple destinations are for future use. There are three possible states:

- Normal—No destination is logging or merging.
- Logging—At least one destination is logging; none are merging.
- Merging—At least one destination is merging.

DR group write mode

Characterizes how a write from a host is replicated. A DR group has two modes of replication writes from the source Vdisk to the destination Vdisk:

- Synchronous replication—I/O requests to both source and destination must be complete before the storage system processes the next I/O request from the host.
- Asynchronous replication—The storage system processes the next I/O request on the source even if there are outstanding I/O requests on the destination. This mode is not available at this time.

DR mode

The operational mode of a DR group that indicates the capability of I/O to be written to its source and/or its destination. There are two options:

- Active at the source
- Passive at the destination

dual fabric

Two independent fabrics providing multipath connections between FC end devices.

EIA

Electronic Industries Alliance. A standards organization specializing in the electrical and functional characteristics of interface equipment.

Enterprise Virtual Array

An HP StorageWorks product that consists of one or more storage systems. Each storage system consists of a pair of HSV controllers and the disk drives they manage. A storage system within the Enterprise Virtual Array (EVA) can be formally referred to as an Enterprise Storage system, or generically referred to as a storage system.

event

Any change that is significant to the storage system. Events include:

- A state change in hardware
- A state change in a logical element, such as a virtual disk
- A completion of a procedure
- An environmental change
- An operational failure
- A change in configuration, such as a new virtual disk that has been created or a new physical disk that has been inserted

fabric

A network of Fibre Channel switches or hubs and other devices.

failover

This term is context specific.

- DR group failover—An operation to reverse the direction of a DR group.
- Managed set failover—An operation to reverse the direction of all the DR groups in the managed set.
- Fabric or path failover—The act of transferring I/O operations from one fabric or path to another.
- Controller failover—The assumption by one controller of the workload of its partner.

failsafe locked

A mode set to cause a condition in which a failsafe-enabled DR group is unable to complete a write to its destination, or when a member Vdisk fails. This condition prohibits any further host I/O to any of the Vdisks in the DR group, and requires immediate intervention.

failsafe mode

A DR group mode in which all source Vdisks in the group become both unreadable and unwritable if any of their corresponding destination Vdisks is unreachable. No logging takes place for the Vdisks. There are two states:

- Enabled
- Disabled

FATA

Fibre Attached Technology Adapted. Near-online disk drives used for data reference, data archival, data replication, and with applications that use the drives infrequently.

FC

See Fibre Channel.

FCA

Fibre Channel adapter. An adapter used to connect the host server to the fabric. Also called a host bus adapter (HBA).

FC connection

A Fibre Channel path between two storage systems or between a host and its storage. A connection is made up of multiple FC links.

FC link

A path between two adjacent FC ports.

Fibre Channel

An ANSI standard name given to a low-level protocol for a type of serial transmission. The Fibre Channel specifications define the physical link, the low-level protocol, and all other pertinent characteristics.

Fibre Channel adapter

See FCA.

Fibre Channel connection

See FC connection.

Fibre Channel link

See FC link.

flush cache

The act of writing data from cache to a storage medium.

full copy

A copy operation in which all 1 MB blocks written on a source Vdisk since it was created are replicated to a destination Vdisk.

GBIC

Gigabit interface converter. Devices inserted into the ports of the Fibre Channel switch that hold the Fibre Channel cables. A GBIC converts fiber optic cable connections to Fibre Channel switch connections.

GUI

Graphical user interface. A software interface that uses a computer's graphic systems to make a program more user friendly.

heterogeneous SAN support

The ability for the product to operate with different operating systems and storage systems in a SAN.

high availability (HA)

Redundant systems, software, and IT processes to reduce the risk of downtime. No single point of failure.

Home

The preferred storage system for the source Vdisks of a DR group. By default, this is the storage system on which a source Vdisk is created, although this designation is user settable.

homogeneous SAN support

Implies the ability for the product to operate with only homogeneous operating systems and homogeneous storage systems in a SAN.

hop

One interswitch link.

host

A computer that runs user applications and uses (or can potentially use) one or more virtual disks created and presented by the controller pair.

host bus adapter (HBA)

See FCA.

HSV

Hierarchical Storage Virtualization. The name given to the virtualization controller architecture.

I/O module

Input/Output module. The enclosure element that is the FC-AL interface to the host or controller. I/O modules are bus speed specific: either 1 Gb or 2 Gb.

IDX

The 2-digit decimal number portion of the HSV controller termination code display that defines one of 32 locations in the termination code array that contains information about a specific event.

in-band communication

Communication that uses the same communications pipe as the operational data. *See also* out-of-band communication.

initialization

A configuration step that binds the controllers together as a storage array and establishes preliminary data structures on the disk array. Initialization also sets up the first disk group, called the *default disk group*, and makes the storage system ready for use.

ISL

Intersite link or interswitch link. The abbreviation is context sensitive.

leveling

Distribution of all user data within the virtual disks across all physical disks within the disk group.

license key

A license key is required to operate the HSV controller software. The license key is a WWN-encoded sequence that is obtained from the HP license key fulfillment website. Three types of license keys exist:

- Basic VCS—Required to unlock VCS. The basic license key covers both controllers in the storage system. The basic license key is required for system initialization.
- Business Copy—Needed to unlock the snapshot and Snapclone features. This license can be added any time **after** the system has been initialized.
- Continuous Access EVA—Needed to implement the data replication features of the HSV controller software. It is also needed to run the Continuous Access user interface.

local site

See primary site.

log

Storage that is used for logging.

logging

Usually context sensitive. In Continuous Access EVA, logging refers to the history of host write commands (and data) when the destination array is not accessible. If failsafe mode is enabled, logging does not occur.

LUN

Logical unit number. An identifier through which a Vdisk is presented to a host.

M-series switches

Fibre Channel Director and Edge switches made by McDATA and sold by HP. Refer to the *HP StorageWorks Continuous Access EVA Design Reference Guide* for a listing of these switches.

managed set

Any collection of DR groups selected by the user for the purpose of managing them. For example, a managed set can be created to manage all DR groups whose sources reside in the same cabinet or all DR groups that deal with a particular set of applications.

master controller

Whichever controller of the controller pair powers up first.

merge or merging

Transferring the contents of the log to the destination Vdisk in order to synchronize the source and destination Vdisks.

near-online storage

On-site storage of data on media that takes only slightly longer to access than online storage kept on high-speed disk drives. The class of devices using this storage typically have lower performance in one or more metrics (for example, random access) than the devices belonging to the online class.

normalization

The initial full copy that occurs between a source Vdisk to the destination Vdisk.

OCP

Operator control panel. The element that displays the controller's status using LEDs and an LCD. Information selection and data entry is controlled by the OCP pushbuttons.

Open SAN Manager

See OSM.

online storage

An allotment of storage space that is available for immediate use, such as a peripheral device that is turned on and connected to a server. The class of devices using this type of storage typically have higher performance than those belonging to the near-online class.

operation state

Current operating condition of a system component. There are three states:

- Normal
- Failed
- Attention (indicates possible problem)

out-of-band communication

Communication that uses a different communications pipe than that used by operational data. *See also* in-band communication.

preferred path

A preference for which controller of the controller pair manages the virtual disk. This preference is set by the user through Command View EVA when creating the virtual disk. A host can change the preferred path of a virtual disk at any time. The primary purpose of preferring a path is load balancing.

presentation

The process whereby a controller presents a virtual disk only to the host computer that has authorized access.

primary site

Controller pair residing at the data center that contains the source data. A storage system may be considered a primary site for some Vdisks and an alternate site for other Vdisks.

reconstruction

The process of regenerating the data contents from a failed drive member. The reconstruction process writes the data to a spare set disk and incorporates the spare set disk into the mirrorset, striped mirrorset or RAID set from which the failed member came.

relationship

The arrangement created when two storage systems are partnered for the purpose of replicating data between them.

remote site

See alternate site.

reservation state

Three possible reservation states exist for a virtual disk:

- None—No host has a reservation on the virtual disk.
- Regular—One host has a reservation on the virtual disk. A regular reservation will not be preserved through failovers or power failures.
- Persistent—One or more hosts have the virtual disk reserved. A persistent reservation is normally preserved through failovers and power failures. A persistent reservation can be released by other hosts.

Resume

Command issued to a DR group or managed set that causes replication to resume after being suspended. This command may initiate a merging of the DR group log or a full copy.

selective presentation

The process whereby a controller presents a virtual disk only to the host computer that is authorized access.

SFP

Small form factor pluggable. A 2-gigabit per second GBIC.

single path

A single connection or path between storage systems containing source and replicated data, or between a host and the storage assigned to that host.

site failover

Command to change a destination role to a source role at the designated site. *See also* failover.

slave

Whichever controller of the pair powers up last. Also called storage system slave.

Snapclone

A virtual disk that can be manipulated while the data is being copied. Only an active member of a virtual disk family can be Snapcloned.

The Snapclone, like a snapshot, reflects the contents of another virtual disk at a particular point-in-time. Unlike a snapshot, the Snapclone is an actual clone of the source virtual disk and immediately becomes an independent active member of its own virtual disk family.

snapshot

A temporary virtual disk that reflects the contents of another logical disk at the particular point-in-time. A snapshot operation is performed only on an active virtual disk. The active disk and its snapshot constitute a virtual family.

source Vdisk

A Vdisk that contains original data that is replicated to its destination Vdisk.

Storage Management Appliance

The HP OpenView Storage Management Appliance is a specialized server on which the Command View EVA, CA user interface, and other SAN applications run.

storage pool

The aggregated blocks of available storage in the total physical disk array.

storage system

A pair of HSV controllers and the array of physical disks they control. A storage system may contain Vdisks that are sources as well as Vdisks that are destinations. Sometimes used in context as follows:

- Source storage system—Used in the context of a particular DR group, this is the storage system at which the source Vdisk resides.
- Destination storage system—Used in the context of a particular DR group, this is the storage system at which the destination Vdisk resides.

storage virtualization

The transparent abstraction of storage at the block level. It separates out logical data access from physical per disk, or per array, data access.

Suspend

Command issued to a DR group or managed set that temporarily halts replication of I/O from source Vdisks to destination Vdisks in that DR group. Source Vdisks continue to run I/O locally, and the I/O is also copied to the DR group log. The command cannot be issued if the DR group is failsafe enabled.

synchronous mode

The mode of operation of a DR group where the data is written to the source and destination caches, after which a completion acknowledgement is sent to the host. *See also* DR group write mode.

throttling I/O

The concentration or channeling of I/O to specific DR groups or managed sets by the use of Suspend and Resume commands.

topology

An interconnection scheme that allows multiple Fibre Channel ports to communicate. Point-to-point, arbitrated loop, and switched fabric are all Fibre Channel topologies.

uninitialized EVA

An uninitialized system cannot be used for storage. When an uninitialized system is initialized, the management agent creates a preliminary structure that makes the system ready for use. An uninitialized system says “uninitialized” at the top of its properties table and has a warning at the bottom.

UUID

Universal Unique Identifier. A unique 128-bit identifier associated with HSV objects.

VCS

Virtual Controller Software. The firmware that runs the storage system.

Vdisk

A simulated disk drive accessible from hosts attached to the SAN. When it is a member of a DR group, a Vdisk is the HSV implementation of a copy set, and it can have two states: *normal* and *copying*.

Virtual Controller Software

See VCS.

virtual disk

See Vdisk.

virtual disk family

A virtual disk and its snapshot, if a snapshot exists, constitute a family. The original virtual disk is called the active disk. When you first create a virtual disk family, the only member is the active disk.

Vraid0

A virtualization technique that provides no data protection. The data host is broken down into chunks and distributed on the disks that make up the disk group from which the virtual disk was created. Reading and writing to a Vraid0 virtual disk is very fast and makes the fullest use of the available storage, but provides no data protection (redundancy) unless there is parity.

Vraid1

A virtualization technique that provides the highest level of data protection. All data blocks are mirrored, or written twice, on separate physical disks. For read requests, the block can be read from either disk, which can increase performance. Mirroring takes the most storage space because twice the storage capacity must be allocated for a given amount of data.

Vraid5

A virtualization technique that uses parity striping to provide moderate data protection. Parity is a data protection mechanism for a striped virtual disk, on which the data to and from the host is broken down into chunks and distributed to the physical disks that make up the disk group in which the virtual disk was created. If the striped virtual disk has parity, another chunk (a parity chunk) is calculated from the set of data chunks and written to the physical disks. If one of the data chunks becomes corrupted, the data can be reconstructed from the parity chunk and the remaining data chunks.

World Wide LUN Name

A 128-bit identifier associated with a Vdisk (64 bits come from the controller WWN).

World Wide Name

A 64- or 128-bit identifier that uniquely identifies the address of a component on the fabric.

WWLN

See World Wide LUN Name.

WWN

See World Wide Name.

Index

A

- adding a DR group to a managed set [82](#)
- alternate site [18](#), [19](#)
- asynchronous mode [36](#)
- audience [13](#)
- authorized reseller, HP [16](#)

B

- backing up the configuration [84](#)
- bidirectional replication [18](#), [33](#)
- block mapping [28](#)
- bootless DR group failover (Linux) [197](#)
- B-series switches
 - restrictions [57](#)
 - zoning [62](#)
- building sequential disk groups [45](#)
- Business Copy
 - concept [31](#)
 - license [69](#), [177](#), [178](#)

C

- codes
 - event [201](#)
 - termination [201](#)
- Command View EVA
 - creating disk groups [69](#)
 - creating host folders [70](#)
 - creating hosts [70](#)
 - creating Vdisk folders [70](#)
 - creating Vdisks [71](#)
 - description [22](#)
 - disk group hardware failure [143](#)

- initialization [69](#)
- monitoring [67](#)
- presenting copy set to destination host [79](#)
- presenting Vdisks to hosts [72](#), [77](#), [78](#)
- recommended uses [66](#)
- using folders [48](#)
- verifying storage allocation [78](#)
- component repair vs. failover [107](#)
- configuration planning [43](#)
- configuration text files [84](#)
- configuring
 - controller-to-switch connections [61](#)
 - dual fabrics [49](#), [50](#)
 - Fibre Channel switches [61](#)
 - hosts [68](#)
 - host-to-switch connections [62](#)
 - single fabrics [50](#)
 - single FCA [50](#), [52](#)
 - Storage Management Appliance [66](#)
- connected system [36](#)
- connections
 - controller-to-switch [61](#)
 - host-to-switch [62](#)
 - SMA-to-switch [65](#)
- Continuous Access EVA
 - concepts [33](#)
 - configuration [18](#), [59](#)
 - failover [103](#)
 - hardware [19](#)
 - host operating systems [24](#)
 - license [25](#), [69](#)
 - load balancing [52](#)
 - overview [17](#)

- platform zoning requirements [53](#)
- prerequisites [13](#)
- related documentation [13](#)
- required software [22](#)
- restrictions [56](#)
- saving storage configuration [89](#)
- single management zone [53](#)
- support procedures [177](#)
- types of failovers [39](#)
- zoning [41](#)
- Continuous Access user interface
 - accessing [72](#)
 - data replication [113](#)
 - description [22](#)
 - icon descriptions [110](#)
 - main window [75](#)
 - managed sets [38](#)
 - monitoring [67](#)
 - recommended uses [66](#)
 - saving storage configuration [89](#)
 - using folders [48](#)
- controller virtualization algorithms [47](#)
- controller-to-switch cabling [62](#)
- conventions
 - document [14](#)
 - text symbols [15](#)
- copy sets [34](#)
 - creating [75](#)
 - deleting [81](#)
 - presentation to destination host [79](#)
- creating
 - copy sets and DR groups [75](#)
 - destination Snapclone before full copy [177](#)
 - disk groups [69](#)
 - host folders [70](#)
 - hosts [70](#)
 - managed sets [81](#)
 - Vdisk folders [70](#)
 - Vdisks [70](#)
- C-series switches
 - restrictions [57](#)
 - zoning [64](#)

D

- data distribution [32](#)
- data migration [32](#)
- data movement using a Snapclone [178](#)
- data recovery [32](#)
- data replication [33](#)
- Data Replication Manager, HSG80 zoning [41](#), [64](#), [99](#)
- deleting
 - copy sets [81](#)
 - managed sets [83](#)
- destination storage system [18](#)
- destination Vdisk [33](#)
- disable failsafe command [115](#)
- disaster tolerance [18](#)
- disk drive capacity [44](#)
- disk drive performance [44](#)
- disk failure protection level [44](#)
- disk group [28](#)
 - configuration planning [46](#)
 - creating [69](#)
 - description [44](#)
 - hardware failure
 - definition [142](#)
 - on source storage system [143](#)
 - on the destination storage system [150](#)
- document conventions [14](#)
- DR group
 - bootless failover (Linux) [197](#)
 - creating [75](#)
 - creating from Snapclone [179](#)
 - depiction [35](#)
 - description [34](#)
 - failover [106](#)
- DR group properties [36](#)
- DR mode [36](#)
- dual-fabric configuration [49](#), [50](#)

E

- editing a managed set [82](#)
- Element Manager for HSG [22](#), [64](#), [99](#)

Enterprise Virtual Array
 description 19
 typical rack configuration 20
environment 41
event codes 201
event monitoring 67
event scenarios 115
expansion panel 19

F

failover 19
 command 114
 concept 39
 controller 39, 106
 defined 106
 DR group 39, 106
 fabric or path 39, 106
 managed set 39, 106
 planned procedure 119
 planned scenario 106, 116
 storage system 106
 unplanned procedure 127
 unplanned scenario 106, 116
 when to 107
failover vs. component repair 107
failsafe mode 36, 105
failsafe-locked 105
FATA drives 44, 79
Fibre Channel adapter
 in host 21, 70
 in SMA 21
 installation 60
 restriction 56
Fibre Channel switch
 configuration 61
 connections to controller 61
 description 20
folders 48
full copy 37, 177
fully allocated snapshot 31

G

getting help 15
gigabit interface converters 20

H

hardware configuration 60
high availability 21, 49
Home storage system 34
homogeneous SANs 53
host folders, creating 70
host operating systems 24
host zoning 63
hosts
 configuration 68
 creating 70
host-to-switch connections 62
HP OpenVMS
 privileges 122, 127, 138
HP resources
 authorized reseller 16
 storage website 16
 technical support 15
HSG80 controllers
 zoning 64

I

initialization 44, 69
inoperative disk group 142
interswitch link 49

J

Java Runtime Environment 22

L

license keys 69
licensing
 Business Copy EVA 25, 69
 Continuous Access EVA 25, 69
load balancing 28, 52
local site 18
log disk

- description [37](#)
- group membership [79](#)
- in DR group [34](#)
- marked for full copy [177](#)
- log states [37](#)
- logging [143](#)
- loop switches [19](#)
- loss of redundancy [142](#)

M

- managed set [38](#)
 - adding a DR group [82](#)
 - creating [81](#)
 - deleting [83](#)
 - editing [82](#)
 - failover [106](#)
 - removing DR group [83](#)
- merging [37](#), [177](#)
- mirroring [29](#)
- M-series switches
 - restrictions [57](#)
 - zoning [64](#)

N

- naming restrictions [48](#)
- normal mode [105](#)
- normal object names [48](#)
- Notification Utility [67](#)

O

- original state [34](#)

P

- parity striping [29](#)
- path failover [106](#)
- physical storage, concept [28](#)
- placement of log into disk group [79](#)
- planned failover [106](#), [119](#)
- planned transfer of operations [120](#)
- planning
 - configurations [43](#)

- disk group configuration [46](#)
- fabric configurations [49](#)
- zones [53](#)
- platform kits [68](#)
- preferred path settings [71](#)
- prerequisites [13](#)
- presenting a copy set to a destination host [79](#)
- presenting Vdisks to hosts [72](#)
- primary site [18](#), [19](#)

R

- redundant fabrics [18](#)
- related documentation [13](#)
- relationship, storage system [34](#)
- remote copy [33](#)
- remote copy feature [33](#)
- remote mirrors [32](#)
- remote site [18](#)
- removing DR group from a managed set [83](#)
- required hardware [19](#)
- restricted characters [48](#)
- restrictions [56](#)
- Resume command [37](#), [114](#), [178](#)
- resumption of operations procedure [132](#)
- resumption of operations scenario [117](#)
- return operations to Home storage system [134](#)
- return operations to Home storage system scenario [117](#)
- return operations to replaced new hardware [117](#), [134](#)
- reversed state [34](#)

S

- saving storage configuration [89](#)
- Secure Path [71](#)
 - description [24](#)
 - installing [72](#)
- Selective Management [64](#)
- single component failure [107](#)
- single FCA configuration [50](#), [52](#)
- single-fabric configurations [50](#)

- site failover, description 107
 - small form factor pluggables 20
 - Snapclone
 - as a Vdisk 34
 - creating a DR group 179
 - creating before full copy 177
 - data movement 178
 - description 32
 - log spare capacity 37
 - snapshots 31
 - SNMP traps 67
 - software configuration 65
 - source storage system 18
 - source Vdisk 33
 - specifying disk group membership for a log 79
 - stale data 143
 - storage array 19
 - Storage Management Appliance
 - activating the standby SMA with HSG80 controllers 99
 - active 19
 - configuration 19
 - considerations for managing storage 88
 - description 21
 - Device Home Page 73
 - enabling management when HSG80 controllers are present 99
 - HSG80 control 64
 - login window 73
 - management zones 54
 - moving storage management to another SMA 93
 - restarting applications 92
 - software Home page 74
 - software installed 22
 - software setup 66
 - stopping applications 90
 - switch connections 65
 - zone restrictions 53
 - storage system failover 106
 - storage system names 48
 - Storage System Scripting Utility 84, 137
 - stranded capacity 46
 - striping 29
 - support procedures 177
 - supported operating systems 24
 - supported switch products 21
 - Suspend command 36, 114, 178
 - symbols in text 15
 - synchronous mode 36
- ## T
- technical support, HP 15
 - text symbols 15
 - throttling I/O 106
 - traditional RAID 30
 - troubleshooting
 - intersite link problems 175
 - storage problems 153
- ## U
- unplanned failover 106, 116
 - unplanned failover procedure 127
 - unplanned transfer of operations 120
- ## V
- Vdisk folders, creating 70
 - Vdisk, creating 70
 - vertical disk groups 46
 - Virtual Controller Software
 - description 22
 - virtual RAID types 29, 30
 - virtual storage, concept 28
 - virtualization concepts 28
 - Vraid0 29
 - Vraid1 29, 37
 - Vraid5 29
- ## W
- websites
 - HP storage 16
 - HP website 15
 - technical support 15

when or when not to failover [108](#)
when to fail over [107](#)
World Wide Name [54](#), [60](#)
write mode [36](#)
WWN [70](#)

Z

zoning
 B-series switches [62](#)

concept [41](#)
C-series switches [64](#)
host [63](#)
input form [55](#)
M-series switches [64](#)
planning [53](#)
SMA management [54](#)
WWN address [63](#)